

Conversations with Search Engines: SERP-based Conversational Response Generation

PENGJIE REN, Shandong University & University of Amsterdam, China & The Netherlands

ZHUMIN CHEN, Shandong University, China

ZHAOCHUN REN, Shandong University, China

EVANGELOS KANOULAS, University of Amsterdam, The Netherlands

CHRISTOF MONZ, University of Amsterdam, The Netherlands

MAARTEN DE RIJKE, University of Amsterdam & Ahold Delhaize, The Netherlands

In this paper, we address the problem of answering complex information needs by conducting *conversations with search engines*, in the sense that users can express their queries in natural language, and directly receive the information they need from a short system response in a conversational manner. Recently, there have been some attempts towards a similar goal, e.g., studies on Conversational Agents (CAs) and Conversational Search (CS). However, they either do not address complex information needs in search scenarios, or they are limited to the development of conceptual frameworks and/or laboratory-based user studies.

We pursue two goals in this paper: (1) the creation of a suitable dataset, the Search as a Conversation (SaaC) dataset, for the development of pipelines for conversations with search engines, and (2) the development of a state-of-the-art pipeline for conversations with search engines, **Conversations with Search Engines (CaSE)**, using this dataset. SaaC is built based on a multi-turn conversational search dataset, where we further employ workers from a crowdsourcing platform to summarize each relevant passage into a short, conversational response. CaSE enhances the state-of-the-art by introducing a supporting token identification module and a prior-aware pointer generator, which enables us to generate more accurate responses.

We carry out experiments to show that CaSE is able to outperform strong baselines. We also conduct extensive analyses on the SaaC dataset to show where there is room for further improvement beyond CaSE. Finally, we release the SaaC dataset and the code for CaSE and all models used for comparison to facilitate future research on this topic.

CCS Concepts: • **Information systems** → **Search interfaces; Question answering; Top-k retrieval in databases.**

Additional Key Words and Phrases: Conversational modeling, Search engine, Dataset, Neural model

ACM Reference Format:

Pengjie Ren, Zhumin Chen, Zhaochun Ren, Evangelos Kanoulas, Christof Monz, and Maarten de Rijke. 2021. Conversations with Search Engines: SERP-based Conversational Response Generation. *ACM Transactions on Information Systems* 1, 1 (May 2021), 28 pages. <https://doi.org/xx.xxxx/xxxxxxx.xxxxxxx>

Zhumin Chen and Zhaochun Ren are corresponding authors.

Authors' addresses: Pengjie Ren, Shandong University & University of Amsterdam, Qingdao & Amsterdam, China & The Netherlands, p.ren@uva.nl; Zhumin Chen, Shandong University, Qingdao, China, chenzhumin@sdu.edu.cn; Zhaochun Ren, Shandong University, Qingdao, China, zhaochun.ren@sdu.edu.cn; Evangelos Kanoulas, University of Amsterdam, Amsterdam, The Netherlands, e.kanoulas@uva.nl; Christof Monz, University of Amsterdam, Amsterdam, The Netherlands, c.monz@uva.nl; Maarten de Rijke, University of Amsterdam & Ahold Delhaize, Amsterdam, The Netherlands, m.derijke@uva.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1046-8188/2021/5-ART \$15.00

<https://doi.org/xx.xxxx/xxxxxxx.xxxxxxx>

1 INTRODUCTION

As we surround ourselves with a range of mobile devices, e.g., smartphones, smartwatches, which only have small screens or even no screen at all, search is increasingly performed in a conversational manner.¹ Despite this development, for complex information needs where a user's intent may be unclear or where it is not obvious what the single best direct answer should be (if any), complex search engine result pages (SERPs) are still the dominant format to present results to users. SERPs are typically characterized by a diverse set of snippets, usually grouped along vertical dimensions and/or by modality, which is far from our natural mode of communication through conversations. Hence, even when we interact with search engines, the more natural mode of interaction, instead

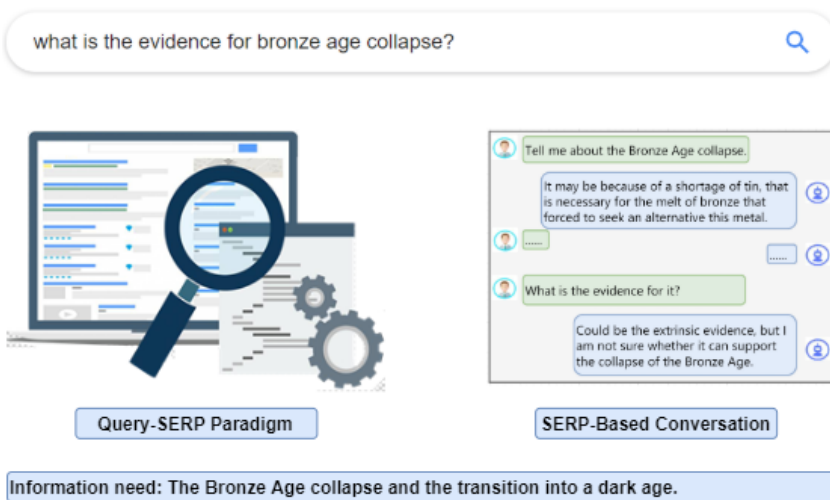


Fig. 1. Search using a traditional SERP vs. a conversation with a search engine, illustrated using the information need *The Bronze Age collapse and the transition into a dark age*.

of a complex SERP is conversational in nature [4, 5, 13, 39]. Figure 1 illustrates the difference given the information need “*The Bronze Age collapse and the transition into a dark age*.” In a *traditional SERP* scenario, we would use keywords to express our information need. For each query, we issue the keywords to a search engine and receive a SERP with a ranked list of results, possibly with snippets, in return. Then, we go through the list and find the information we need from the relevant snippets and/or results. If not, we reformulate our query and this process is repeated until our information need is satisfied. Alternatively, we can fulfill our information need through a *conversation with the search engine*. To achieve this, there are several aspects to be considered, e.g., asking clarifying questions [2, 67], conversational search and recommendation [11, 14, 50, 58], chitchats [61], user intent understanding [38], feedback analysis [25], conversation management [54], and conversational question answering [41], etc.

In this work, we study conversational question answering in search scenarios, i.e., SERP-based conversational response generation, in the sense that we would express our need in natural language and we would directly receive the information we need in a short system response which is a summary of relevant results listed in the SERP in a conversational manner. Most studies on conversational question answering are based on knowledge bases (a.k.a question answering over knowledge base (KBQA)) [65] or a background passage (a.k.a machine reading comprehension (MRC)) [10, 40]. In these scenarios, users have well-specified and specific information needs and

¹<https://www.thinkwithgoogle.com/consumer-insights/personal-needs-search-trends/>

their queries are mostly factoid questions that can be answered by a relatively short text span (entity mentions, etc.) extracted from the given background knowledge (knowledge base, documents, etc.). Several previous studies address complex user queries, e.g., on the QA over knowledge bases task. However, the complex queries in these studies refer to compositional questions that involve joining multiple relations [44]. They are still factoid questions that can be answered by entity mentions or triples from knowledge bases. In contrast, search engines cater for much broader needs/queries, which results in far more complex user queries and search scenarios. Besides, despite the fact that some complex information needs can be answered by extracting a short span from SERPs in a single interaction, many require a deep understanding of multiple SERPs and multiple interactions are needed in order to achieve all the information needs, e.g., the one shown in Figure 1. There are some studies towards conversational question answering in search scenarios, but they all have critical limitations. For instance, the CAsT dataset² only provides ground truth passages as answers and does not have conversational responses. And the MS MARCO question answering (QA) dataset [33] is single-turn and when there is no answer, it just leaves the system response blank, which is sufficient for training models but not suitable for evaluation. Hence, these datasets are not sufficient to support the development of conversations with search engines.

To address this gap we pursue two goals: (1) the creation of a suitable dataset for the development of pipelines for SERP-based conversational response generation, and (2) the development of state-of-the-art baseline components that make up such a pipeline using this dataset. The dataset, called *Search as a Conversation* (SaaC), is developed in a Wizard-of-Oz fashion [6]. We simulate users based on conversational queries from the CAsT dataset. Then, we employ online workers (a.k.a., “wizards”) to play the role of the system. The wizards have access to SERPs from which they can get useful information to respond to the user queries. We ask the wizards to find supporting sentences from results/snippets on a SERP and summarize them into short conversational responses. When there is no direct answer, we ask the wizards to generate something that is likely relevant, e.g., “It could be the extrinsic evidence, but I am not sure ...”, or interesting to the user, e.g., “I have no idea about how melatonin was discovered. But I can tell you that ...”.

As to our second main goal in this paper, using the SaaC dataset for the development of pipelines to support conversations with search engines, we devise a modularized multi-task learning framework, called *Conversations with Search Engines* (CaSE). CaSE decomposes conversations with search engines into four sub-tasks: (1) conversation & passage understanding (CPU), (2) relevant passage selection (RPS), (3) supporting token identification (STI), and (4) response generation (RG). CPU is a module aiming at understanding and encoding conversations and passages. RPS then finds relevant passages based on the encoded representations from the CPU. STI further identifies supporting tokens that are eventually used in the response. Finally, RG generates the response based on the output from the above three modules.

Because there are no ground truth labels for STI to define a supervised learning loss, we present a weakly-supervised Confidence-Critical Cross Entropy (CCCE) learning loss based on the intuition that the overlapping tokens between the ground truth responses and the passages are more likely supporting tokens than the non-overlapping ones, especially overlapping larger rare tokens. In order to make CaSE generate more accurate responses, we propose a Prior-aware Pointer Generator (PPG) to implement RG by considering the passage and token probabilities from RPS and STI as priors so that the generated responses are expected to be more accurate by including supporting tokens from relevant passages.

²<http://www.treccast.ai/>

We conduct experiments to: (1) compare the performance of state-of-the-art methods from related tasks to our CaSE, (2) understand the contribution of the four sub-tasks in CaSE, and (3) identify room for further improvement on SaaC beyond CaSE.

To sum up, the contributions of this work are as follows:

- We introduce the task of conversations with search engines and build a new dataset, the SaaC dataset.
- We decompose the task into four sub-tasks (CPU, RPS, STI, RG), and propose a modularized CaSE model that uses a weakly-supervised CCCE loss to identify supporting tokens and PPG to encourage generating more accurate responses.
- We conduct extensive experiments to show the effectiveness of CaSE and identify room for further improvements on conversations with search engines.

In Section 2 we describe the SaaC dataset and its creation. In Section 3 we describe the CaSE model. Section 4 details our experimental setup and the outcomes of our experiments are presented in Section 5. Related work is discussed in Section 6 and we conclude in Section 7.

2 THE SaaC DATASET

Our goal is to create a dataset containing interactions between the users and a search engine where they communicate through conversations to achieve search tasks. Instead of creating a new dataset from scratch, we build such a dataset based on the TREC CASt dataset. In this section we describe the stages involved in creating the SaaC dataset.

2.1 Collecting Conversational Queries

TREC CASt [14] has already built a collection of conversational query sequences, so we reuse their data to reduce development cost. Here, we briefly recall the process used in collecting the TREC CASt data.

The topics in the CASt dataset are collected from a combination of previous TREC topics (Common Core [9], Session Track [23], etc.), MS MARCO Conversational Sessions, and the CASt organizers [14]. The organizers ensured that the information needs are complex (requiring multiple rounds of elaboration), diverse (across different information categories), open-domain (not requiring expert domain knowledge to access), and mostly answerable (with sufficient coverage in the passage collection). A description of an example topic is shown in Figure 2 (top half). Then, the TREC CASt organizers created sequences of conversational queries for each turn. They started with the description of the topic and manually formulated the first conversational query. After that, they formulated follow-up conversational queries by introducing coherent transitions, e.g., coreference and omission. For example, “Tell me about the Bronze Age collapse. ... What is the evidence for it? (introducing coreference).”, or “What is a physician’s assistant? ... What’s the average starting salary in the UK? ... What about in the US? (introducing omission).” There is a constraint that later conversational queries only depend on the previous queries, but not on system responses. We will discuss this constraint later. The reader is referred to [14] for a more detailed account of the creation of the TREC CASt dataset.

2.2 Collecting Candidate Passages

For our dataset creation, we combine three standard TREC collections: MARCO Ranking passages, Wikipedia (TREC CAR), and News (Washington Post) as the passage collection. To introduce more complex passages and meanwhile achieve higher recall for the current query, we follow Voskarides et al. [57] to extend the current query by extracting words that capture relevant information from previous turns and add them to the query of the current turn. Next, we use standard query likelihood

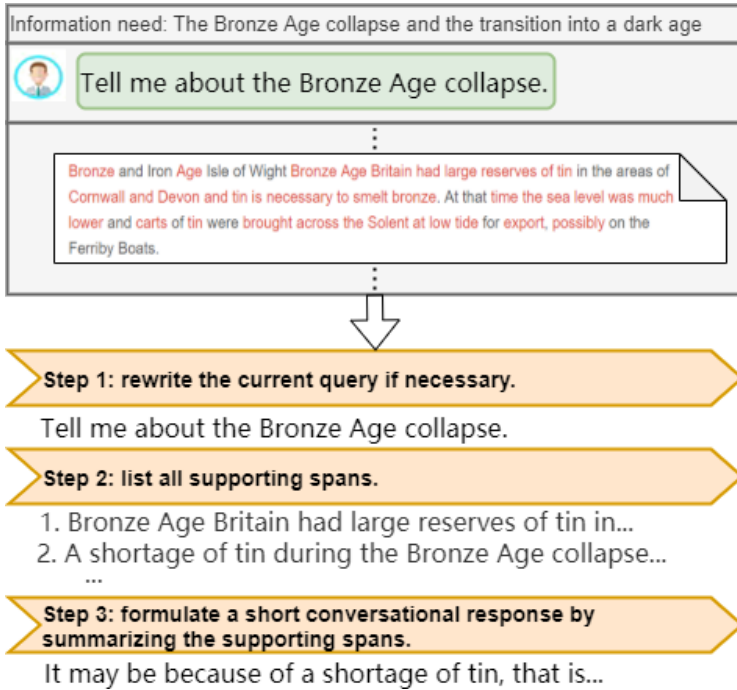


Fig. 2. The process of building the SaaC dataset.

Table 1. Comparison of conversational datasets on search, chitchat and question answering.

Dataset	Multiple Natural language turns	Multiple query	Multiple Abstractive sources	Multiple response
Holl-e [32], WoW [16]	✓	✓	✗	✓
QuAC [10], CoQA [41]	✓	✗	✗	✗
CAsT [14]	✓	✓	✓	✗
MS MARCO [33]	✗	?	✓	✓
SaaC (this paper)	✓	✓	✓	✓

with Dirichlet smoothing and RM3 relevance feedback as the ranking model to retrieve the top 10 candidate passages (if the ground truth passage is within the top 10, otherwise, we retrieve the top 9 and manually add the ground truth passage). Note that although we rewrote the current query to make it self-contained (which will be detailed in the next subsection), we did not use the rewritten queries when preparing the candidate passages in order to stay close to practical search engines. Finally, we randomly shuffle the top 10 candidate passages to eliminate position bias.

2.3 Collecting Conversational Responses

We employ online workers from Amazon Mechanical Turk (MTurk)³ to collect conversational responses in a Wizard-of-Oz fashion, where we ask the workers to play the role of the system and write responses based on the provided passages, as shown in Figure 2. Specifically, we first rewrite the queries if necessary and we require that the rewritten queries should be self-contained (step 1). Then, the workers need to list all supporting spans from the passages that contain facts to help

³<https://www.mturk.com/>

generate the responses (step 2). The supporting spans are kept the same as they are in the passages. Finally, we ask the workers to summarize the supporting spans into short, conversational responses. Other requirements include: (1) make sure the responses are case sensitive and grammatically correct; (2) avoid using the spans directly in the responses without summarization; and (3) talk about secondary relevant information or information that could be interesting to the users, when no answer can be found in the passages. We take three actions to guarantee that the collected data meets our requirements: First, we only employ high-quality master workers of MTurk who are required to continue to pass the statistical monitoring to retain the qualification. Second, we provide detailed guidelines, which also include several good and bad examples with detailed explanations. Third, we manually checked each annotation by ourselves to guarantee the agreement between different workers, as it is hard to automatically check the quality of these responses written in natural language. We make sure they all meet the given requirements: (1) The queries are answered directly with further explanations; (2) The responses contain everything that is needed to answer the queries; (3) The supporting spans are summarized instead of being used directly in the responses; (4) The responses are grammatically and syntactically correct, and case sensitive.

In total, 28 master MTurk workers participated in our task. The number of annotations from each worker ranges from 1 to 244. Table 1 shows that while SaaC shares some characteristics with existing datasets, it also has its own unique characteristics, tailored to support research on conversations with search engines. For example, Holl-e and WoW contain multi-turn conversations, however they target chitchat. QuAC and CoQA have multi-turn conversations, but the passages are all Wikipedia articles and fixed for all conversation turns; the responses are extractive text spans from Wikipedia articles. The datasets listed above do not target search scenarios, where the passages are from various sources, e.g., blogs, Encyclopaedia, social media, and there are queries that cannot be answered properly with extractive text spans. CAsT is for search scenarios, but it does not have abstractive responses. Out of all the datasets available at the time of writing, MS MARCO is the closest to SaaC. However, MS MARCO is single-turn and not all queries are in natural language. Table 2 lists some instances from both datasets to show their differences.

2.4 Description of the SaaC dataset

The SaaC dataset has 80 topics (with a total of 748 queries, and 7–12 queries per topic) from CAsT. Almost all of the queries come with a complex information need; 359 are “what” queries, 144 are “how” queries, and 47 are “why” queries. It is important to note that, compared with CAsT, SaaC has three types of additional annotations: the reformulated queries, the supporting spans from the relevant passages, and the conversational responses.

We report some descriptive statistics for the SaaC dataset in Table 3; for comparison we include the same information for the MS MARCO dataset. In the table, “#pairwise passage similarity” denotes the average TF-IDF dot similarity of each pair of candidate passages. “#n-gram overlap” denotes the average n-gram overlap ratio between the answer and candidate passages. “query/answer common words ratio” denotes the ratio of common words (word frequency $\geq 100,000$) in the query and answer, respectively. We can see from Table 3 that SaaC is more complex and conversational than MS MARCO. SaaC is more complex in that (1) the “#query/answer/passage length” is larger, which means the queries are more complex to understand, and the passages contain more (noisy) information; and (2) the “#pairwise passage similarity” is larger, which means it has more confusing candidate passages, making it hard to find the correct ones. SaaC is also more conversational because: (1) the “n-gram overlap” between the answers and passages is much smaller, which means the answers are more abstractive; and (2) the “#query/answer common words ratio” is larger, which means that the queries and answers are more in spoken language, which is more informal.

Table 2. Some examples from the SaaC dataset. For comparison, we also include examples from the MS MARCO dataset.

	Queries	Responses
SaaC	What are the different forms of energy?	As far as I know, the different forms of energy are kinetic energy, potential energy, gravitational energy, nuclear energy, ...
	How can it be stored?	Energy can be stored in many ways like a gallon of gasoline or a barrel of oil contains stored energy that can be released when we burn it.
	What type of energy is used in motion?	The energy of motion is kinetic energy.
	Tell me about mechanical energy.	Mechanical energy can be either kinetic energy (energy of motion) or potential energy (stored energy of position).
	Give me some examples.	Examples of mechanical energy are the power of a football flying through the air or the wrecking ball that is swung backward away from a building.
	Tell me about the Bronze Age collapse.	It may be because of a shortage of tin, that is necessary for the melt of bronze that forced to seek an alternative of this metal.
	What is the evidence for it?	Could be the extrinsic evidence, but I am not sure whether it can support the collapse of the Bronze Age.
	What are some of the possible causes?	One of the possible causes of the Bronze Age collapse is the invasion of the Sea Peoples.
MS MARCO	Who were the Sea Peoples?	The Sea Peoples was a confederacy of seafaring raiders that caused political unrest, and attempted to enter or control Egyptian territory during the late 19th dynasty, and the 20th dynasty.
	What was their role in it?	The Sea Peoples entered and invaded eastern Mediterranean territory. The Sea People's invasions ushered or caused the bronze age collapse.
	cell organelles definition	Cell organelles is a membrane bound compartment or structure in a cell that performs a special function.
	What music style developed from Scott Joplin?	Scott Joplin developed the maple leaf rag and the entertainer.
	glioma meaning	The meaning of glioma is a tumor springing from the neuroglia or connective tissue of the brain, spinal cord, or other portions of the nervous system.

3 A METHOD FOR CONVERSING WITH SEARCH ENGINES

Users could have conversations with search engines for many reasons, like getting information, performing tasks, etc. In this work, we only focus on users' information exploration needs, in a setting where users express their needs in natural language and the system returns the information the users need in short conversational responses.

Table 3. Descriptive statistics for the SaaC dataset. For comparison, we include the same information for the MS MARCO dataset.

	SaaC	MS MARCO
#query length	7.21	6.08
#answer length	28.19	15.90
#passage length	155.25	67.11
#pairwise passage similarity	31.14%	26.77%
#1-gram overlap	80.75%	90.93%
#2-gram overlap	47.68%	72.92%
#3-gram overlap	34.85%	63.43%
#4-gram overlap	27.85%	58.01%
#query common words ratio	70.80%	56.15%
#answer common words ratio	60.21%	53.61%

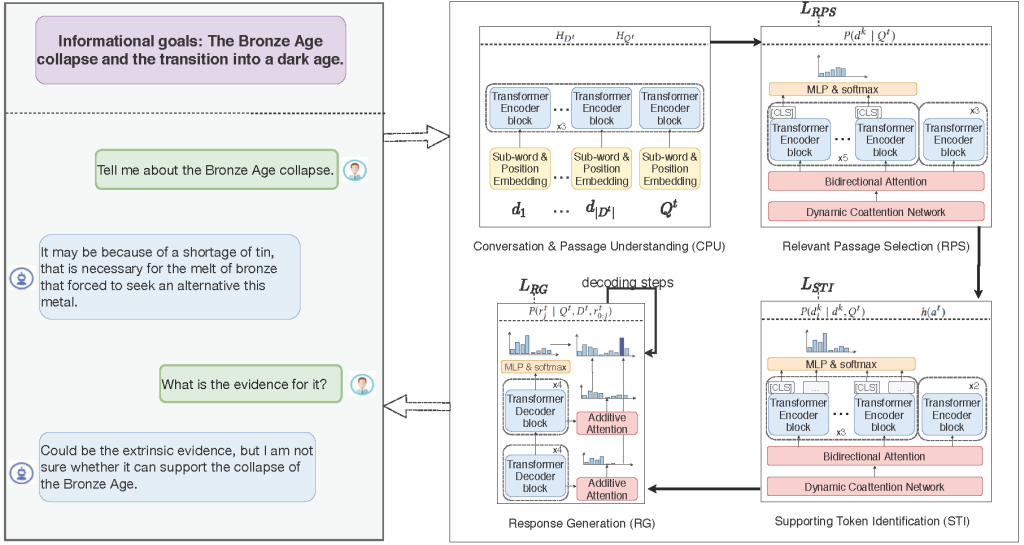
Table 4. Summary of main symbols and notation used in the paper.

q	A user query. If there are superscript and/or subscript, superscript means the conversation turn and subscript means the token index, e.g. q_i^t .
Q^t	A list of queries until t th turn.
d	A passage. If there are superscript and/or subscript, superscript means the passage index in the candidate set and subscript means the token index, e.g. d_i^k .
D^t	The passage candidate list for the t th turn.
h	Hidden representation for a token. The superscript and/or subscript indicate the definition (what it models), the type (query or document), the token index and so on, e.g. $h_{q_i^t}$.
H	A sequence of representations.
r	A response. If there are superscript and/or subscript, superscript means the conversation turn and subscript means the token index, e.g. r_i^t .
$P()$	Probability.
W or v	Learnable parameters.
\hat{y} or y	(Weak) ground truth label.

Table 4 summarizes the main symbols and notation used in this paper. Formally, given a series of previous user queries $Q^{t-1} = [q^{t-m}, q^{t-m+1}, \dots, q^{t-1}]$ in natural language, the user query at the current turn q^t , and a list of candidate passages $D^t = [d^1, d^2, \dots, d^k]$ that have been retrieved by a search engine and potentially contain the answers (or are at least relevant to q^t),⁴ the SaaC task is to generate a short response r^t for q^t by identifying the relevant passages and summarizing the supporting spans from these passages into a conversational response. Note that we do not take the previous system responses as input, because: (1) this might result in error accumulation when the previous system responses are improper in practice; and (2) in our datasets, all the background information that is needed to understand the current query is covered by the previous user queries. But, it is easy to extend all the models including the baselines and ours to consider the previous system responses.

In this work, we decompose the SaaC task into four sub-tasks: (1) conversation & passage understanding (CPU), (2) relevant passage selection (RPS), (3) supporting token identification (STI),

⁴While creating effective queries from multi-turn questions is a challenging task in itself, we assume that the candidate documents are provided in advance so as to simplify the experimental design and facilitate reproducibility.



Conversations with Search Engines (CaSE) system interface

Conversations with Search Engines (CaSE) model architecture

Fig. 3. An overview of Conversations with Search Engines (CaSE). Section 3 contains a walkthrough of the model.

and (4) response generation (RG), as shown in Figure 3. We devise a modularized framework, Conversations with Search Engines (CaSE), to operationalize the sub-tasks in an end-to-end manner. Both MS MARCO and our SaaC still adopt a “user asks – system responds” paradigm which does not consider mixed-initiative [67], so we preclude this part in the CaSE framework.

Specifically, the CPU module first encodes each query q and passage d into a sequence of hidden representations, i.e.,

$$H_{q^t} = [\mathbf{h}_{q_1^t}, \dots, \mathbf{h}_{q_{|q^t|}^t}] \text{ for } q^t, \text{ and} \quad (1)$$

$$H_{d^k} = [\mathbf{h}_{d_1^k}, \dots, \mathbf{h}_{d_{|d^k|}^k}] \text{ for } d^k. \quad (2)$$

Then, based on the query and passage representations, the RPS module selects the relevant passage by estimating the passage relevance probability $P(d^k | Q^t)$ for each passage d^k in the candidate pool D^t . After that, the STI module estimates the probability of each passage token to be a supporting token, $P(d_i^k | d^k, Q^t)$, where a token is supporting if it contributes to the final responses. Finally, the RG module takes the outputs from the above three modules into consideration and generates a short response, token by token. In particular, $P(d^k | Q^t)$ and $P(d_i^k | d^k, Q^t)$ are modeled as priors in the RG module.

In the following subsections, we will describe our proposed solution for each module in detail.

3.1 Conversation and Passage Understanding

We employ a Transformer model to perform conversation and passage understanding, which relies on self-attention to extract important information to represent conversations and passages. Specifically, for the conversational queries $Q^t = Q^{t-1} \cup \{q^t\}$, we first concatenate the tokens as one sequence, and then input it into a stack of Transformer encoder blocks [56] to obtain representations for each query and each token $H_{Q^t} = [H_{q^{t-m}}, H_{q^{t-m+1}}, \dots, H_{q^t}]$, where each H_{q^t} is defined as in Eq. 1. Note that we put a special token “[CLS]” at the start, formally referred to as $\mathbf{h}_{q_{CLS}^t}$, which

is considered to represent the conversations up to the current conversation turn t . Similarly, we obtain representations H_{d^k} as in Eq. 2 for each passage.

3.2 Relevant Passage Selection

In order to model the relevance to conversational queries of each passage, we first need to model the interaction between them. Here, we employ a similar bi-directional attention flow as proposed by Seo et al. [46] to do machine reading comprehension (MRC), which is also used by Nishida et al. [34] to do question answering (QA). Specifically, we first obtain the interaction matrix $M^{Qd_k} \in \mathbb{R}^{|Q^t| \times |d^k| \times N}$ between the conversation tokens H_{Q^t} and each passage H_{d^k} [46] as follows:

$$M^{Qd_k} = \begin{bmatrix} f^{Qd_k}(\mathbf{h}_{q_1^{t-m}}, \mathbf{h}_{d_1^k}) & f^{Qd_k}(\mathbf{h}_{q_1^{t-m}}, \mathbf{h}_{d_2^k}) & \cdots & f^{Qd_k}(\mathbf{h}_{q_1^{t-m}}, \mathbf{h}_{d_{|d^k|}^k}) \\ f^{Qd_k}(\mathbf{h}_{q_2^{t-m}}, \mathbf{h}_{d_1^k}) & f^{Qd_k}(\mathbf{h}_{q_2^{t-m}}, \mathbf{h}_{d_2^k}) & \cdots & f^{Qd_k}(\mathbf{h}_{q_2^{t-m}}, \mathbf{h}_{d_{|d^k|}^k}) \\ \cdots & \cdots & \ddots & \cdots \\ f^{Qd_k}(\mathbf{h}_{q_{|q^t|}^t}, \mathbf{h}_{d_1^k}) & f^{Qd_k}(\mathbf{h}_{q_{|q^t|}^t}, \mathbf{h}_{d_2^k}) & \cdots & f^{Qd_k}(\mathbf{h}_{q_{|q^t|}^t}, \mathbf{h}_{d_{|d^k|}^k}) \end{bmatrix}, \quad (3)$$

where f^{Qd_k} is the cross-correlation function, which captures the correlation between each pair of tokens in queries and passages:

$$f^{Qd_k}(\mathbf{h}_{q_i}, \mathbf{h}_{d_j^k}) = \mathbf{v}^{Qd_k \top} [\mathbf{h}_{q_i} \oplus \mathbf{h}_{d_j^k} \oplus (\mathbf{h}_{q_i} \odot \mathbf{h}_{d_j^k})], \quad (4)$$

where $\mathbf{v}^{Qd_k} \in \mathbb{R}^{3N \times 1}$ is the parameter vector; \oplus denotes the concatenation operation and \odot denotes the Hadamard product. Then, a Dynamic Coattention Network as proposed in [62] is used to obtain the dual attention representations for the conversational queries $H_{Q^t \leftarrow D^t}$, and each passage $H_{d^k \leftarrow Q^t}$ as follows:

$$\begin{aligned} H_{Q^t \leftarrow D^t} &= [H_{Q^t} \oplus H_{Q^t \leftarrow D^t}^1 \oplus H_{Q^t \leftarrow D^t}^2 \oplus (H_{Q^t \leftarrow D^t}^1 \odot H_{Q^t}) \oplus (H_{Q^t \leftarrow D^t}^2 \odot H_{Q^t})] \\ H_{d^k \leftarrow Q^t} &= [H_{d^k} \oplus H_{d^k \leftarrow Q^t}^1 \oplus H_{d^k \leftarrow Q^t}^2 \oplus (H_{d^k \leftarrow Q^t}^1 \odot H_{d^k}) \oplus (H_{d^k \leftarrow Q^t}^2 \odot H_{d^k})]. \end{aligned} \quad (5)$$

Take $H_{Q^t \leftarrow D^t}$, for instance. It contains five parts: H_{Q^t} is the preserved information from the queries; $H_{Q^t \leftarrow D^t}^1$ and $H_{Q^t \leftarrow D^t}^2$ represent the preserved information from the passages after the first- and second-order selection with the help of the interaction matrix M^{Qd_k} ; $(H_{Q^t \leftarrow D^t}^1 \odot H_{Q^t})$ and $(H_{Q^t \leftarrow D^t}^2 \odot H_{Q^t})$ capture the further interaction between two types of information. The same is true for $H_{d^k \leftarrow Q^t}$. Here, $H_{Q^t \leftarrow D^t}^1$, $H_{Q^t \leftarrow D^t}^2$, $H_{d^k \leftarrow Q^t}^1$ and $H_{d^k \leftarrow Q^t}^2$ are defined as follows:

$$H_{Q^t \leftarrow D^t}^1 = \max(\{M_{d^k} H_{d^k}\}_{k=1, \dots, |D^t|}) \quad (6)$$

$$H_{Q^t \leftarrow D^t}^2 = \max(\{M_{d^k} H_{d^k}^1\}_{k=1, \dots, |D^t|}) \quad (7)$$

$$H_{d^k \leftarrow Q^t}^1 = M_Q^\top H_{Q^t} \quad (8)$$

$$H_{d^k \leftarrow Q^t}^2 = M_Q^\top H_{Q^t}^1 \quad (9)$$

$$M_{d^k} = \text{softmax}_d(M^{Qd_k}) \quad (10)$$

$$M_Q = \text{softmax}_Q(M^{Qd_k}). \quad (11)$$

In Equation 6–11, \max denotes max pooling; softmax_Q and softmax_d denote the softmax over M^{Qd_k} along the query and passage dimension, respectively. Then, a stack of the Transformer encoder blocks are used to reduce the dimension of $H_{Q^t \leftarrow D^t}$ and $H_{d^k \leftarrow Q^t}$. By conducting the above operations, the model can learn to control the relevant information flow from queries to passages (e.g., $D^t \leftarrow Q^t$, $d^k \leftarrow Q^t$), and vice versa.

To estimate the passage relevance score, we use an MLP to get the passage relevance score by taking the first token representation of each passage $H_{d_0^k \leftarrow Q^t}$ (corresponding to the “[CLS]” token in Section 3.1) as input. The relevance score is normalized with a sigmoid to obtain the relevance probability $P(d^k | Q^t)$. We use binary cross entropy to supervise the learning of this module:

$$L_{RPS} = \sum_{d^k \in D^t} y_{d^k} \log P(d^k | Q^t) + (1 - y_{d^k}) \log(1 - P(d^k | Q^t)), \quad (12)$$

where $y_{d^k} = 1$ if d is, relevant otherwise $y_{d^k} = 0$.

3.3 Supporting Token Identification

Previous methods directly generate the responses after passage selection, which we hypothesize can be improved by incorporating a dedicated supporting token identification (STI) module. The core idea is that, besides learning to select the relevant passage, the model could also learn to identify supporting tokens, which might be used to generate the response.

To do so, we use a similar architecture as in the RPS module (parameters are not shared) to get updated representations $H_{Q^t \leftarrow D^t}$ and $H_{D^t \leftarrow Q^t}$. But, instead of estimating the passage relevance score based on the “[CLS]” representation (which represents the general representation of the passage), we estimate the probability of each passage token as a supporting token $P(d_i^k | d^k, Q^t)$. Specifically, we use an MLP to get the supporting token likelihood score for each token with $h_{d_i^k \leftarrow Q^t} \in H_{D^t \leftarrow Q^t}$ as input ($h_{d_i^k \leftarrow Q^t}$ represents the token representation for d_i^k), which is normalized with a sigmoid to obtain $P(d_i^k | d^k, Q^t)$.

Unfortunately, there are no ground truth labels to define a supervised learning loss to train $P(d_i^k | d^k, Q^t)$. To this end, we design a weak supervision signal based on the following intuitions: (1) if a passage token d_i^k is a supporting token, it must exist in the ground truth response; (2) if the surrounding tokens of d_i^k are also in the ground truth response, d_i^k is more likely to be a supporting token; and (3) rare passage tokens that exist in the ground truth response are more likely to be supporting tokens than frequent ones. Specifically, we devise a *Confidence-Critical Cross Entropy* (CCCE) loss as follows:

$$L_{STI} = - \sum_{d^k \in D^t} \sum_{d_i^k \in d^k} \left[c(d_i^k) \hat{y}_{d_i^k} \log P(d_i^k | d^k, Q^t) + (1 - \hat{y}_{d_i^k}) \log(1 - P(d_i^k | d^k, Q^t)) \right], \quad (13)$$

where $\hat{y}_{d_i^k}$ is a weak label indicating whether d_i^k is a supporting token; $\hat{y}_{d_i^k} = 1$ if $d_i^k \in r^{t*}$, and 0 otherwise, where r^{t*} is the ground truth response; and $c(d_i^k)$ is a coefficient indicating the confidence of d_i^k as a supporting token, which is defined as:

$$c(d_i^k) \propto \frac{1}{\log \text{freq}(d_i^k)} \cdot \prod_n |d_{i-n:i+n}^k \cap r^{t*}|, \quad (14)$$

where $\text{freq}(d_i^k)$ is the token frequency in the data collection. The first term models how “rare” d_i^k is, the second term models how likely d_i^k and its n surrounding tokens are supporting tokens (overlapping with the ground truth). Finally, $c(d_i^k)$ ensures that rare and more often overlapping tokens get more opportunities to be identified as supporting tokens.

3.4 Response Generation

We propose a Prior-aware Pointer Generator (PPG) to implement the RG module, which is able to generate tokens from a predefined vocabulary and copy tokens from both the queries and passages. The idea is that each token can be generated under three modes, i.e., with probability $P(r_j^t | Q^t, D^t, r_{0:j}^{t*}, g)$ from the generation mode g (a.k.a. the *vocabulary generator* which generates

from a predefined vocabulary), with probability $P(r_j^t | Q^t, D^t, r_{0:j}^t, c_{Q^t})$ from the query copying mode c_{Q^t} (a.k.a. the *query pointer generator* which copies from queries), and with probability $P(r_j^t | Q^t, D^t, r_{0:j}^t, c_{D^t})$ from the passage copying mode c_{D^t} (a.k.a. the *prior-aware passage pointer generator* which copies from passages). The final probability of generating the token r_j^t is a combination of the three probabilities. Especially when estimating the passage copying probability, PPG models the passage relevance and supporting token likelihood from the RPS and STI modules as priors.

Given the previous decoded tokens $r_{0:j}^t = [r_0^t, \dots, r_{j-1}^t]$ (r_0^t is set to a special token “[BOS]” indicating the beginning of decoding), we first use a stack of Transformer decoder blocks [56] to obtain the hidden representations

$$H_{r_{0:j}^t}^Q = [h_{r_0^t}^Q, \dots, h_{r_{j-1}^t}^Q],$$

which takes $r_{0:j}^t$ and $H_{Q^t \leftarrow D^t}$ as inputs. Then, we use the same stack of Transformer Decoder blocks to obtain hidden representations

$$H_{r_{0:j}^t}^D = [h_{r_0^t}^D, \dots, h_{r_{j-1}^t}^D], \quad (15)$$

which takes $H_{r_{0:j}^t}^Q$ and $H_{D^t \leftarrow Q^t}$ as inputs. Afterwards, we estimate the token probability from three modes: generating from the vocabulary g , copying from queries c_{Q^t} , and copying from passages c_{D^t} .

Vocabulary generator. The probability of generating a token from the predefined vocabulary is estimated as:

$$P(r_j^t | Q^t, D^t, r_{0:j}^t, g) = P(g | r_{0:j}^t) \text{softmax}(\text{mlp}([r_{j-1}^t \oplus h_{r_{j-1}^t}^D \oplus h(a^t)])), \quad (16)$$

where $P(g | r_{0:j}^t)$ denotes the probability of the generation mode g ; $h(a^t)$ is the answer representation from the STI module, which is estimated as follows:

$$h(a^t) = \sum_{d^k \in D^t} P(d^k | Q^t) \sum_{d_i^k \in d^k} P(d_i^k | d^k, Q^t) h_{d_i^k \leftarrow Q^t}, \quad (17)$$

where $P(d^k | Q^t)$ is the passage relevance probability; $h_{d_i^k \leftarrow Q^t}$ is the i -th token representation from $H_{d^k \leftarrow Q^t}$. Both are from the RPS module. $P(d_i^k | d^k, Q^t)$ is the supporting token probability from the STI module.

Query pointer generator. We use another additive attention to estimate the probability of copying a token from the conversational queries:

$$P(r_j^t = Q_i^t | Q^t, D^t, r_{0:j}^t, c_{Q^t}) = P(c_{Q^t} | r_{0:j}^t) P(Q_i^t | Q^t, D^t, r_{0:j}^t), \quad (18)$$

where $P(c_{Q^t} | r_{0:j}^t)$ is the query copying mode probability; and $P(Q_i^t | Q^t, D^t, r_{0:j}^t) = \text{attention}(\text{query} : h_{r_{j-1}^t}^Q, \text{key} : H_{Q^t \leftarrow D^t})$.

Prior-aware passage pointer generator. The probability of copying a token from the passages is

$$\begin{aligned} & P(r_j^t = d_i^k | Q^t, D^t, r_{0:j}^t, c_{D^t}) \\ &= P(c_{D^t} | r_{0:j}^t) \sum_{d^k \in D^t} P(d^k | Q^t) \cdot \sum_{d_i^k \in d^k} P(d_i^k | d^k, Q^t) P(d_i^k | Q^t, D^t, r_{0:j}^t), \end{aligned} \quad (19)$$

where $P(c_{D^t} | r_{0:j}^t)$ is the passage copying mode probability; $P(d^k | Q^t)$ is the passage prior from the RPS module and $P(d_i^k | d^k, Q^t)$ is the supporting token prior from the STI module; and $P(d_i^k | Q^t, D^t, r_{0:j}^t) = \text{attention}(\text{query} : h_{r_{j-1}^t}^D, \text{key} : H_{D^t \leftarrow Q^t})$. Here, the prior-aware passage pointer generator learns to copy tokens from all the candidate passages despite that we can predict which

passage is most relevant from the RPS module. The reason is that it is hard for the RPS module to always rank the most relevant passage as top one.

To coordinate the probabilities from different modes, we learn a mode coordination probability:

$$[P(g | r_{0:j}^t), P(c_{Q^t} | r_{0:j}^t), P(c_{D^t} | r_{0:j}^t)] = W^\top [h_{r_{j-1}^t} \oplus h_{Q^t}^{att} \oplus h_{D^t}^{att}], \quad (20)$$

where $W \in 3N \times 3$ is the parameter matrix; $h_{Q^t}^{att}$ and $h_{D^t}^{att}$ are the attended query and passage representations from the two attentions in the query and passage pointer generators, respectively. The final probability at the j -th decoding step is

$$P(r_j^t | Q^t, D^t, r_{0:j}^t) = P(r_j^t | Q^t, D^t, r_{0:j}^t, g) + P(r_j^t | Q^t, D^t, r_{0:j}^t, c_{Q^t}) + P(r_j^t | Q^t, D^t, r_{0:j}^t, c_{D^t}). \quad (21)$$

If a token is absent from a mode, its corresponding probability from that mode is set to zero.

We use the negative log likelihood loss to train PPG as follows:

$$L_{RG} = - \sum_{r_j^t \in r^t} \log P(r_j^t | Q^t, D^t, r_{0:j}^t). \quad (22)$$

4 EXPERIMENTAL SETUP

We seek to answer the following research questions by designing experiments:

- (Q1) What is the performance of CaSE compared to other methods? Does CaSE outperform the state-of-the-art methods in terms of response generation and passage ranking performance? What, if any, are the performance differences on the MS MARCO and SaaC datasets?
- (Q2) What is the effect of different components of CaSE on its overall performance?
- (Q3) Where does CaSE fail? That is, is there any room for further improvement on the SaaC dataset?

4.1 Datasets and Evaluation Metrics

As pointed out in Section 2.4, SaaC contains 80 topics (with 748 queries) from CASt, which is too small to train neural models in an end-to-end fashion. Hence, we train all the models that we consider on the MS MARCO 2.1 Q&A + Natural Language Generation training set (MS MARCO train) [33]. Although this is a single-turn dataset and the queries and responses are less conversational than in the SaaC dataset (see Table 2), it does have passage relevance labels and human written answers, which is sufficient as a training set. MS MARCO is the sole dataset with real queries from search engines and human written answers. There are two types of answer in the MS MARCO dataset: QA style (corresponding to the “answers” field in the dataset) and NLG style (corresponding to the “wellFormedAnswers” field in the dataset). The QA style answers are not always in natural language, according to the official description.⁵ The NLG style answers were generated by rewriting the QA style answers to make sure that: (1) each NLG style answer includes proper grammar to make it a full sentence; (2) the NLG style answers make sense without the context of either the query or the passage; and (3) the NLG style answers have a high overlap with exact portions in one of the context passages. This ensures that NLG style answers are true natural language and not just selected spans. In this work, we focus on generating natural language responses, so we only keep data samples where the “wellFormedAnswers” field is not empty. We randomly split the original development set into two sets with roughly equal size, one as our development set (MS MARCO dev) and the other as our test set (MS MARCO test). The sizes of the MS MARCO train, MS MARCO dev, MS MARCO test, and SaaC test are 524,105, 32,345, 32,225, and 1,008,⁶ respectively.

⁵<https://github.com/microsoft/MSMARCO-Question-Answering>

⁶Some queries have more than one relevant passage. We construct a test sample for each relevant passage.

Although there are other multi-turn conversational datasets, e.g., those shown in Table 1, we do not report results on them as they are not for the search scenario.

We use BLEU (up to 4-grams using uniform weights), and ROUGE-1, ROUGE-2, and ROUGE-L to evaluate the response generation performance, which are commonly used in natural language generation tasks, e.g., QA, MRC [24, 34]. We also report MAP, Recall@5 and NDCG for passage ranking performance.

4.2 Methods Used for Comparison

We collect and implement state-of-the-art methods from various related tasks.

- **S2SA**. S2SA is a sequence-to-sequence with attention model; it is commonly used as a baseline for natural language generation tasks [31, 32].
- **GTTP** [45]. GTTP improves S2SA by incorporating a pointer mechanism that enables it to copy tokens from the input during generation. GTTP achieves state-of-the-art performance on many natural language generation tasks [35, 41].
- **TMemNet** [16]. TMemNet was first introduced for the knowledge grounded dialogue task. It combines a Memory Network and Transformer to do knowledge selection and dialogue response generation.
- **GLKS** [42]. GLKS is a state-of-the-art method for Background-Based Conversations (BBC) (the best performing method on the Holl-E⁷ dataset at the time of writing). It introduces a mechanism to combine global and local knowledge selection for dialogue response generation.
- **Masque** [34]. Masque is the best performing method on the MS MARCO Q&A + Natural Language Generation task at the time of writing. Because we only use “wellFormedAnswers” to train the models, we remove the answer possibility classifier and style token.
- **CaSE**. CaSE is proposed in this paper.

4.3 Implementation Details

For a fair comparison, we implement all models used in our experiments based on the same code framework to ensure that they share the same code apart from the model itself. We set the word embedding size and hidden size to 256. We use the BERT vocabulary⁸ for all methods but for a fair comparison, we avoid using any extra resources for all methods, including pre-trained embeddings. The vocabulary size is 30,522. The learning rate was increased linearly from zero to 2.5×10^{-4} in the first 6,000 steps and then annealed to 0 by using a cosine schedule. We use gradient clipping with a maximum gradient norm of 1. We use the Adam optimizer ($\alpha = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$). An exponential moving average was applied to all trainable variables with a decay rate of 0.9995. For all models, we combine multiple losses linearly if there is more than one. For TMemNet, Masque and CaSE, the parameters of the Transformer encoder for the queries and passages are shared, as we found that this resulted in better performance. We train all models on four TITAN X (Pascal) GPUs. The batch size is chosen from (32, 64, 128) according to the GPU memory. We select the best models based on performance on the development set.

5 EXPERIMENTAL RESULTS

5.1 How Does CaSE Perform?

To answer Q1, we report the results of all methods on both MS MARCO (see Table 5) and SaaC (see Table 6). S2SA, GTTP and GLKS do not perform passage ranking, so there are no MAP, Recall@5 and NDCG results for those methods. Given the results, there are three main lessons.

⁷<https://github.com/nikitaacs16/Holl-E>

⁸<https://github.com/huggingface/transformers>

Table 5. Overall performance on the MS MARCO test set (%). **Bold face** indicates the best result in terms of the corresponding metric. R-1: ROUGE-1; R-2: ROUGE-2; R-L: ROUGE-L. * and ** indicate CaSE is significantly better than Masque (p-value < 0.05 and p-value < 0.01 with t-test, respectively). Note that the scores are not comparable to the leaderboard of MS MARCO at <http://www.msmarco.org/leaders.aspx> as different training and test sets are used.

Methods	MS MARCO test						
	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
S2SA	46.75	25.17	38.79	18.61	–	–	–
GTTP	47.41	25.74	39.64	19.32	–	–	–
GLKS	47.74	27.51	40.01	20.75	–	–	–
TMemNet	48.01	29.56	40.29	23.23	53.43	82.56	64.83
Masque	55.33	37.15	47.75	30.01	63.75	90.77	72.82
CaSE	57.44**	39.26**	49.91**	32.08**	65.75**	92.20**	74.34**

Table 6. Overall performance on the SaaC test set (%). **Bold face** indicates the best result in terms of the corresponding metric. R-1: ROUGE-1; R-2: ROUGE-2; R-L: ROUGE-L. * and ** indicate CaSE is significantly better than Masque (p-value < 0.05 and p-value < 0.01 with t-test, respectively).

Methods	SaaC test						
	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
S2SA	24.71	09.04	17.69	07.23	–	–	–
GTTP	28.89	10.51	20.61	08.80	–	–	–
GLKS	35.46	13.46	25.31	12.38	–	–	–
TMemNet	35.29	12.72	25.05	09.60	14.12	11.31	29.25
Masque	35.44	13.48	25.74	11.97	16.13	12.00	31.77
CaSE	37.34*	14.58	27.24*	13.26	17.22	13.72	32.67

First, CaSE achieves the best results on both datasets in terms of all metrics. Especially, CaSE outperforms the best performing model Masque on the MS MARCO leaderboard and the best performing model GLKS from the BBC task (at the time of writing). Generally, CaSE improves over Masque by around 2%pt in terms of generation evaluation metrics and around 1–2%pt in terms of passage ranking metrics. Part of the improvement is from the proposed STI and PPG modules, which we will analyze in more detail in §5.2.

Second, the scores on the MS MARCO dataset are much higher than those on the SaaC dataset. This is the case for all methods, including CaSE. For example, the BLEU score of CaSE is 18.82%pt higher on MS MARCO for response generation, and the MAP score is 48.53%pt higher. This confirms that SaaC is a more challenging and more suitable dataset than MS MARCO for research on conversations with search engines. SaaC is more challenging because (1) the queries and passages are more complex, which is clear from Table 3. (2) SaaC has greater query, passage and answer lengths and the passages are more similar. SaaC is more suitable for conversations with search engines because (1) it has multi-turn conversations, which introduces the requirements of modeling context from historical turns. (2) the responses are more abstractive and conversational, which is closer to real conversation scenarios. See Section 5.3 for further details.

Third, modeling passage selection is necessary. On the one hand, we can see that the methods with passage selection modules (TMemNet, Masque and CaSE) are generally much better than those

Table 7. Ablation study on the MS MARCO test set (%). CaSE-X denotes CaSE with the component X left out. CaSE-RG: replacing the RG component with a traditional pointer generator. CaSE-STI: CaSE without the STI component. CaSE-RPS: CaSE without the RPS module. There is no CaSE-CPU because all other components rely on the CPU component.

CaSE variants	MS MARCO test						
	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
CaSE	57.44	39.26	49.91	32.08	65.75	92.20	74.34
CaSE-RG	57.21	38.91	49.62	31.58	65.16	91.95	73.90
CaSE-STI	55.68	36.94	47.92	29.98	64.02	91.10	73.03
CaSE-RPS	56.36	38.37	48.75	31.22	–	–	–

Table 8. Ablation study on the SaaC test set (%). CaSE-X denotes CaSE with the component X left out. CaSE-RG: replacing the RG component with a traditional pointer generator. CaSE-STI: CaSE without the STI component. CaSE-RPS: CaSE without the RPS module. There is no CaSE-CPU because all other components rely on the CPU component.

CaSE variants	SaaC test						
	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
CaSE	37.34	14.58	27.24	13.26	17.22	13.72	32.67
CaSE-RG	38.32	15.20	28.21	14.45	16.54	12.47	32.51
CaSE-STI	37.74	14.51	27.62	13.01	16.70	12.68	32.39
CaSE-RPS	37.83	15.02	26.87	14.39	–	–	–

without (i.e., S2SA, GTTP and GLKS). On the other hand, we also notice that the improvements for passage ranking are consistent with the improvements for response generation.

5.2 What Do the Components of CaSE Contribute?

To answer Q2 and analyze the effects of the RPS, STI and RG modules in CaSE, we conduct an ablation study. We do not separately study the modeling and learning of the STI module, as removing the L_{STI} loss will result in a lack of direct supervision to guide the learning of STI, which does not make sense. The results on the MS MARCO and SaaC test sets are shown in Table 7 and 8, respectively.

Generally, removing any module will result in a drop in performance in terms of both response generation and relevant passage selection. Specifically, the results drop by more than 2%pt in terms of BLEU and more than 1.5%pt in terms of MAP on the MS MARCO dataset by removing the STI module. This is also true for the RPS and RG modules, although the drops are not as large as for STI. The generation performance of CaSE-STI is even worse than Masque, which confirms the effectiveness of STI. Interestingly, although we found that modeling passage selection is very important (Table 5 and 6) for the other models, removing the RPS module does not much influence the overall performance of CaSE. We think the reason is that CaSE incorporates the STI module which has some overlapping effects with RPS to some extent, as when a passage contains more supporting tokens, it is more relevant in general. To sum up, even if the RPS, STI and RG share some common effects, they are also complementary to each other as combining them will bring further improvements.

One exception is that CaSE-RG achieves better response generation performance than CaSE on SaaC. We believe that the reason for this behavior is that although CaSE can generate more accurate

responses by leveraging the outputs from the RPS and STI modules as priors in PPG (which can be verified by the better performance on MS MARCO), this will influence the abstractiveness of the response, because CaSE is encouraged to put more emphasis on the tokens in the relevant passages with PPG and directly copy them to the responses instead of generating them from the vocabulary. As a result, CaSE performs better on the MS MARCO dataset while CaSE-RG performs better on the SaaC dataset, because SaaC is more abstractive, as can be seen from Table 3.

5.3 Is there Room for further Improvement?

To answer Q3, we explore the room for further performance improvements on the SaaC dataset by conducting additional experiments and/or case studies.

Table 9. Response generation and passage ranking performance of CaSE on the SaaC dataset w.r.t. different conversational turns (%).

Turn	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
1	47.97	28.13	38.22	33.16	0.2117	0.1909	0.3521
2	38.05	16.27	27.00	15.56	0.1607	0.1043	0.3205
3	42.11	17.25	30.54	13.95	0.1904	0.1289	0.3541
4	37.84	13.61	26.08	13.45	0.1635	0.1234	0.3127
5	36.28	11.80	26.74	10.39	0.1406	0.0944	0.2992
6	31.73	08.33	21.92	04.53	0.1361	0.0930	0.2950
7	30.42	08.32	21.95	05.04	0.1451	0.1120	0.3000
8	33.05	11.03	24.39	08.46	0.1060	0.0936	0.2415
≥ 9	35.16	13.88	26.06	11.00	0.1672	0.1088	0.3445

Table 10. Response generation and passage ranking performance of Masque on the SaaC dataset w.r.t. different conversational turns (%).

Turn	R-1	R-2	R-L	BLEU	MAP	Recall@5	NDCG
1	47.42	27.02	36.55	28.49	0.2039	0.1550	0.3485
2	30.99	11.82	21.70	11.52	0.1670	0.1182	0.3174
3	34.94	13.38	24.40	09.41	0.2162	0.1607	0.3689
4	39.87	17.51	29.34	17.41	0.1536	0.1114	0.3039
5	35.85	10.53	27.20	06.40	0.1628	0.0977	0.3272
6	34.04	11.51	24.00	11.42	0.1479	0.1423	0.3115
7	31.26	09.86	22.55	08.05	0.1449	0.1155	0.3143
8	29.84	06.95	20.04	05.12	0.0754	0.0339	0.2158
≥ 9	32.38	10.02	23.13	07.27	0.1624	0.1279	0.3330

First, more effort is needed to properly model the context, i.e., the conversational history, of the current turn query. To illustrate this, we show the performance of response generation and passage ranking of each turn in Table 9 and 10. We found that CaSE is better than Masque in general on the SaaC dataset mostly because CaSE performs better for the first three turns, and the improvements come mainly from the generation part. This actually means that the proposed STI and RPS modules do indeed help to generate better responses, as without the two modules CaSE basically degenerates into Masque. However, it is worth to note that the advantage of CaSE fades away as the conversational turns go on, which means that the main bottleneck becomes how to better model the context.

Table 11. Response generation performance of CaSE on the SaaC dataset (%). OQ: Using the original query of the current turn; CQ+OQ: Using context queries + original query of the current turn; RQ: Using reformulated ground truth query of the current turn; GP: Using ground truth passages; GS: Using ground truth supporting sentences.

	R-1	R-2	R-L	BLEU
OQ	39.24	18.71	29.49	18.05
CQ+OQ	37.34	14.58	27.24	13.26
RQ	42.79	22.17	32.75	22.05
GP	46.19	25.78	33.37	27.05
GS	48.74	29.39	36.54	32.34

We see that the performance is much higher in the first turn because there is no context that needs to be considered. Performance drops dramatically for the following turns including the 2-nd turn, and performance tends to get worse as the number of turns increases. There is an exception for the ≥ 9 -th turns, which are better than the 8-th turn in terms of passage ranking. This may be because hard queries do not go beyond the 8-th turn in CASt. We analyzed the queries from the ≥ 9 -th turns and found that there are a lot of “what” queries in these two turns like “What type does chemical energy belong to?”, which generally needs less modeling of missing context. A deeper understanding of the current query is challenging because it is not just a matter of coreference resolution [1]. It is common that people omit information to keep the conversations natural, which is well reflected in the SaaC dataset.

Second, more effort is needed to obtain a better understanding of the current turn query. To illustrate this, we conduct a comparison of using the original current query (OQ), the context queries and the current original query (CQ+OQ), and the reformulated current query (RQ), as shown in Table 11. RQ is a manually rewritten version of the current original query, where we make sure that RQ is self-contained and does not need to rely on context. We can see that using RQ improves the performance, which is not surprising. But we should also note that even when using RQ, the performance is far from perfect as it does not even outperform the 1st turn in Table 9. This indicates that the model cannot create a understanding of the current query even though all needed information is provided. This is confirmed by the fact that CQ+OQ performs worse than OQ: the model is doing worse by involving more contextual information.

Third, more effort is needed to perform better at passage selection and support identification. We illustrate this by showing the results of CaSE when the ground truth passages (GP) or sentences (GS) are provided in Table 11. We can see that although CaSE has used effective and complex mechanisms to perform relevant passage selection and supporting token identification, there is still a lot of room to improve in this direction.

Fourth, more effort is needed to investigate how to generate more conversational and abstractive responses. Due to limitations of the MS MARCO data, the models are rarely trained to learn to generate tokens that address the conversational nature of responses. For instance, given the query “What is arnica used for?,” the current model will just list the answers, “arnica , trauma , pain and shock.” The human response is “Arnica is a plant based remedy used to relief pain. It is also used to speed injury and trauma healing as well as to reduce bruising.” Besides, there are some cases where there is no answer or only a partial answer is available in the passages; the current model will either generate a wrong answer or just leave it blank. However, in practical scenarios, the system should indicate it does not know the answer or only knows part of the answer and reply accordingly, e.g., “Sorry, I don’t know much about the largest tiger shark ever to have lived on Earth or caught. But I do know the largest great white shark ever caught on camera, it was a seven

meter-long female, called Deep Blue.” We can build suitable datasets to address this. Alternatively, we can investigate how to leverage datasets from related tasks, e.g., chitchat datasets, where there are more natural and conversational human responses [16, 32].

5.4 Case study

In order to obtain an understanding of what the responses look like, in this subsection we perform a qualitative analysis by examining some generated responses from CaSE and Masque. We present 5 examples from the MS MARCO test set as shown in Table 12. We also show 2 conversation examples from the SaaC test set as shown in Table 13 and 14. We use the following ratings to indicate the quality of each response based on the ground truth responses. ++: very good (The generated responses are fluent, coherent, and contain all relevant information as in ground truth responses); +: good (The generated responses are mostly fluent, coherent, and contain most relevant information as in ground truth responses); 0: fair (The generated responses are readable, and contain some relevant information as in ground truth responses); -: bad (The generated responses are readable but contain wrong information); --: very bad (The generated responses are not readable at all or contain irrelevant information only).

From Table 12, we can see that both CaSE and Masque perform well in general on the MS MARCO test set. In most cases, they can find the most relevant passages, based on which they can generate proper responses. The generated responses are fluent and coherent with the queries. For example, when it is a ‘why’ question, both models generate responses starting with ‘because’. We can also see that CaSE is indeed better than Masque generally. This is mostly because that CaSE is better at selecting the more relevant passages. For instance, in Example 4 and 5, CaSE generates perfect responses by identifying the most relevant passages and supporting tokens while Masque generates improper or wrong responses because of the selection of the irrelevant passages. This indicates that it is effective that CaSE uses the proposed STI module and its corresponding loss to help better select the relevant passage and further identify the supporting tokens. However, there are also cases when both models select the secondary relevant or irrelevant passages instead of the most relevant ones. For instance, in Example 3, both models generate the same response which tries to explain the meaning of glioma. Although the generated response is acceptable, it is not as good as the human-written response where more aspects about glioma are covered.

Table 13 shows a conversation example on the topic “forms of energy” where CaSE performs well. In most turns, CaSE can generate proper responses which means that CaSE can understand the queries and select the relevant passages. Especially, compared with Masque, CaSE is better at handling queries at later turns. For instance, in Turn 3, it seems that Masque is confused with the queries in Turn 2 and 3 and the generated response is a mixture of the answer to both Turn 2 and 3. As on the MS MARCO test set, CaSE performs better in terms of selecting the relevant passages and identifying the supporting tokens. This is illustrated in Turn 4, where CaSE selects the most relevant passage and generates the same response as the human response by identifying the right supporting tokens. In contrast, the response generated by Masque is acceptable but not as good as CaSE. The main reason is still that Masque selects the secondary relevant passage.

It should be noted that there are still some cases where both CaSE and Masque perform badly. An example on the topic “Bronze Age collapse” is shown in Table 14. We can see that, for most turns, both CaSE and Masque generate unsatisfactory responses. The reasons for the poor performance are three-fold. First, the topics are more difficult in that the passage is longer and there are more confusing concepts. For instance, for Turn 1, there are 71 passages annotated as relevant but only 3 of them are annotated as most relevant. A passage is annotated as most relevant only if it contains all the relevant information to form a complete answer. The most relevant passages contains the complete answers As a result, it is difficult for the models to find the passage that contains the

Table 12. Some examples from the MS MARCO test set.

Example 1	Query	cell organelles definition	
	Human	cell organelles is a membrane bound compartment or structure in a cell that performs a special function .	
	CaSE	a membrane - bound compartment or structure in a cell that performs a special function .	+
	Masque	cell organelles are a membrane - bound compartment or structure in a cell that performs a special function .	++
Example 2	Query	what music style developed from scott joplin ?	
	Human	scott joplin developed the maple leaf rag and the entertainer .	
	CaSE	the maple leaf rag and the entertainer .	+
	Masque	the music of ragtime , scott joplin was the foremost composer of the genre in the early 20th century , known for works like the maple leaf rag and the entertainer .	+
Example 3	Query	glioma meaning	
	Human	the meaning of glioma is a tumor springing from the neuroglia or connective tissue of the brain , spinal cord , or other portions of the nervous system .	
	CaSE	a glioma is a type of tumor that starts in the brain or spine .	0
	Masque	a glioma is a type of tumor that starts in the brain or spine .	0
Example 4	Query	why does steam hammer occur ?	
	Human	it occur by the accumulation of condensate (water) trapped in a portion of horizontal steam piping .	
	CaSE	because of the accumulation of condensate (water) trapped in a portion of horizontal steam piping .	++
	Masque	because it is caused by a steam bubble forming or being pushed into a pipe completely filled with water .	--
Example 5	Query	how wide is a two way road ?	
	Human	a two way road is between 44 and 48 feet wide .	
	CaSE	a two way road is between 44 and 48 feet wide .	++
	Masque	a two way road is 50 or 60 feet wide .	-

answers from a couple of passages which are similar but do not contain the answers or only contain part of the answers. This can be further confirmed by some of the statistics, e.g., ‘#passage length’, ‘#pairwise passage similarity’ in Table 3. Second, the follow-up queries depend more on the previous queries and modeling of long-term dependency might be needed. For comparison, for the example in Table 14, in order to understand the queries in later turns (Turn 4 and 5), the approaches still need an understanding of the background ‘Bronze Age collapse’ from the the query. Third, there are cases where no candidate passage contains a direct answer but just some secondary relevant information instead. For example, we found that, in the SaaC dataset, there are many human-written responses containing statements like “could be”, “I think”, “I am not sure”, or even “I don’t know”, which indicate that when the secondary relevant passages are provided, even for humans, it is sometimes not easy to write the responses. Finally, the above analysis also indicates that the SaaC dataset brings new challenges that remain to be solved in the future.

Table 13. An example for the topic “forms of energy” from the SaaC test set.

Turn 1	Query	what are the different forms of energy ?	
	Human	as far as i know , the different forms of energy are kinetic energy , potential energy , gravitational energy , nuclear energy , chemical energy , heat energy , electrical energy , electromagnetic energy , sound energy and solar energy .	
	CaSE	light energy , heat energy , mechanical energy , gravitational energy , electrical energy , sound energy , chemical energy , nuclear or atomic energy .	+
	Masque	light energy , heat energy , mechanical energy , gravitational energy , electrical energy , sound energy , chemical energy , nuclear or atomic energy .	+
Turn 2	Query	how can it be stored ?	
	Human	energy can be stored in many ways like a gallon of gasoline or a barrel of oil contains stored energy that can be released when we burn it .	
	CaSE	the different forms of energy can be stored in the guelph ultra store .	-
	Masque	the different forms of energy can be stored in the form of glycogen in the liver and muscles , to be available for later use .	-
Turn 3	Query	what type of energy is used in motion ?	
	Human	the energy of motion is kinetic energy .	
	CaSE	kinetic energy is the energy of motion .	++
	Masque	the different forms of energy is stored in the kinetic energy .	--
Turn 4	Query	tell me about mechanical energy .	
	Human	mechanical energy can be either kinetic energy (energy of motion) or potential energy (stored energy of position) .	
	CaSE	mechanical energy can be either kinetic energy (energy of motion) or potential energy (stored energy of position) .	++
	Masque	mechanical energy is the energy which is possessed by an object due to its motion or due to its position .	0
Turn 5	Query	give me some examples .	
	Human	examples of mechanical energy are the power of a football flying through the air or the wrecking ball that is swung backward away from a building .	
	CaSE	an example of mechanical energy is the power of a football flying through the air .	+
	Masque	mechanical energy is the sum of kinetic and potential energy in an object that is used to do work .	--

6 RELATED WORK

We briefly present an overview of related work on Conversational Search (CS) and on Conversational Agents (CAs).

6.1 Conversational Search

The concept of search as a conversation has been around since the 1980s [4, 13]. Until recently, the idea did not attract a lot of attention due to limitations in data and computing resources at the time. Now, the topic is back in the spotlight. One branch of work conducts user studies on CS. For example,

Table 14. An example for the topic “Bronze Age collapse” from the SaaC test set.

Turn 1	Query	tell me about the bronze age collapse .	
	Human	it may be because of a shortage of tin , that is necessary for the melt of bronze that forced to seek an alternative this metal .	
	CaSE	robert drews describes the collapse as " the worst disaster in ancient history , even more calamitous than the collapse of the western roman empire . "	--
	Masque	robert drews describes the collapse as " as 230 years , corresponding to the uncertainty regarding the length of the " dark age " of mesopotamia with regard to the uncertainty regarding the length of the " .	-
Turn 2	Query	what is the evidence for it ?	
	Human	could be the extrinsic evidence , but i am not sure whether it can support the collapse of the bronze age .	
	CaSE	the bronze age collapse .	--
	Masque	the bronze age collapse and trade disruptions in the mediterranean around 1300 bc , forced metalworkers to seek an alternative to bronze .	--
Turn 3	Query	what are some of the possible causes ?	
	Human	one of the possible causes of the bronze age collapse is the invasion of the sea peoples .	
	CaSE	the iron age in europe is being seen as a part of the bronze age collapse in the ancient near east , in ancient india (with the post - rigvedic vedic civilization) , ancient iran ,	--
	Masque	the greek tradition relates this destruction to the dorians , but it is suggested that the dorian invasion was only one of the causes of the bronze age collapse in the eastern mediterranean , as there is no evidence	+
Turn 4	Query	who were the sea peoples ?	
	Human	the sea peoples was a confederacy of seafaring raiders that caused political unrest , and attempted to enter or control egyptian territory during the late 19th dynasty , and the 20th dynasty .	
	CaSE	the bronze age collapse is the transition from the late bronze age to the early iron age , expressed by the collapse of palace economies of the aegean and anatolia , which were replaced after a hiatus by the isolated	--
	Masque	the sea peoples was a confederacy of seafaring raiders who sailed into the eastern shores of the mediterranean , caused political unrest , and attempted to enter or control egyptian territory during the late 19th dynasty , and	+
Turn 5	Query	what was their role in it ?	
	Human	the sea peoples entered and invaded eastern mediterranean territory . the sea people ' s invasions ushered or caused the bronze age collapse .	
	CaSE	the sea peoples ushered the bronze age collapse , which resulted in the cultural collapse of mycenaean greece , the hittite empire , the new kingdom of egypt and the civilizations of canaan and syria .	+
	Masque	robert drews describes the collapse as " " " " " " " " .	--

Vtyurina et al. [59] conduct a user study, where they ask 21 participants to solve 3 information seeking tasks by conversing with three agents: an existing commercial system, a human expert, and a perceived experimental automatic system, backed by a human “wizard behind the curtain.” They

show that existing conversational assistants cannot be effectively used for complex information search tasks. Vakulenko et al. [55] argue that existing studies neglect exploratory search when users are unfamiliar with the domain of their goal. They investigate interactive storytelling as a tool to enable exploratory search within the framework of a conversational interface. Trippas et al. [53] conduct a laboratory-based observational study for CS, where pairs of people perform search tasks communicating verbally. They conclude that CS is more complex and interactive than traditional search.

Another line of work has proposed theoretical frameworks concerning CS. Radlinski and Craswell [39] present a theoretical framework of information interaction in a chat setting for CS, which highlights the need for multi-turn interactions. Azzopardi et al. [3] propose a conceptual framework that outlines the actions and intents of users and agents in order to enable the user to explore the search space and resolve their information need. The work listed above studies CS either in a theoretical or a user study environment. The theoretical/conceptual frameworks have made requirements about the data annotations more demanding, often going beyond currently available datasets.

Zhang et al. [68] devise a System Ask-User Respond paradigm for CS, and design a memory network for product search and recommendation in e-commerce. Lei et al. [27] initiate a path breaking Estimation—Action—Reflection framework to achieve deep interaction between recommendation and conversation [60]. Following this line, Aliannejadi et al. [2] and Zamani et al. [67] formulate the task of asking clarifying questions in information retrieval. Bi et al. [7] propose a conversational paradigm for product search, and an aspect-value likelihood model to incorporate both positive and negative feedback on non-relevant items. Zamani and Craswell [66] introduce an open-source framework with a modular architecture for CS. It has a modular design to encourage the study of new CS algorithms. It can also integrate with a user interface to facilitate user studies. To advance research on CS, a number of CS datasets have been released recently.⁹ For example, Gustavo Penha and Hauff [19] release a multi-domain dialogue dataset containing information-seeking conversations from the Stack Exchange community.¹⁰ Dalton et al. [14] organize a TREC Conversational Assistance Track (CAST), which establishes a concrete and standard collection of data with information needs to make systems directly comparable. In the first year, the benchmark only focuses on candidate passage retrieval: Read the dialogue context and perform retrieval over a large collection of passages.

Although the studies listed above propose concrete datasets or methods for CS, none of them targets directly generating conversational system responses by modeling CS in a search engine scenario. In contrast, instead of just retrieving relevant passages, we propose to directly generate conversational responses for CS, which is more challenging as we need to not only find the relevant passage but also summarize it into a short conversational response. Note that although summarization is involved in this process, it is quite different from query-driven summarization [43], where: (1) the queries are more like a general information need description, e.g., the query “Describe steps taken and worldwide reaction prior to introduction of the Euro on January 1, 1999. Include predictions and expectations reported in the press.” from DUC 2007;¹¹ (2) the passages are mostly from articles like news reports which are in formal language; and (3) the answer is usually much longer, e.g., on DUC 2007, answers can be as long as 250 words. Recently, Thomas et al. [52] and Trippas et al. [54] create the MISC dataset and the CSCData dataset for spoken conversational search by recording pairs of seeker and intermediary collaborating on information seeking. However, there

⁹<https://github.com/chauff/conversationallIR>

¹⁰<https://stackexchange.com/>

¹¹<https://duc.nist.gov/duc2007/tasks.html>

are no supporting documents available and the responses are in very verbose spoken language. In most cases, the responses are more like a description of SERPs instead of a direct summary of the answers.

6.2 Conversational Agents

Conversational modeling has long been a hot research topic in natural language processing [21, 29]. Most research falls into three groups: task-oriented dialogue system (TDS) agents [47, 48], social bots, and QA agents [17]. TDS aims to achieve a specific task for users, e.g., booking a flight [8, 28], while social bots aim to satisfy the human need for communication and so on [16, 31]. These goals are quite different from people’s goals in search scenarios where user information needs can be more complex and exploratory.

Efforts to build QA agents come in two main flavors: KBQA and MRC, which study how to query a KB interactively with natural language, and generate an answer to an input query based on a set of passages, respectively. Existing methods are either retrieval/extraction-based [15, 20] or generation-based [65]. The former try to query an entity from a KB [64], retrieve evidence from evidence candidates [26], or extract a span from a given passage as direct answers [40]. The latter generates natural language answers directly [51]. Recently, QA has been extended to multi-turn conversational scenarios [10], which introduce more challenges related to conversational understanding. For example, Christmann et al. [12] target complex questions in KB-QA, which involve joining multiple relations and multi-hop reasoning in a conversational setting [18, 49]. Qu et al. [36] study how to retrieve evidence from a large collection in a multi-turn conversational context. On many benchmark datasets, the best models are approaching human performance [10, 44] or even surpass humans [41]. However, the studies listed above mostly target fact-centric information needs that can be answered by entity mentions or triples from a knowledge base (KB) [22]. Besides, the passages are fixed for all conversation turns and from a single source, i.e., Wikipedia, which is different from search scenarios where SERPs are generated with information from multiple sources. Some approaches target complex questions that cannot be well answered by KB triples [24, 37]. In particular, Nguyen et al. [33] collect a large-scale dataset, MS MARCO, from Bing usage logs, where the answers are written by real humans to ensure that they are in natural language.

The work listed above focuses on single-turn QA. In this work, we extend it to multi-turn conversational scenarios where query understanding, relevant passage finding and response generation, etc. are more challenging. In order to achieve better relevant passage finding and response generation, we further propose the CaSE model, which improves a state-of-the-art model by introducing the STI module and a weakly-supervised CCCE loss.

7 CONCLUSION AND FUTURE WORK

In this paper, we propose *conversations with search engines* as task for the community to consider and we contribute two types of result: First, we release a new test set, SaaC, which is more suitable and challenging for this research than existing resources. Second, we propose an end-to-end neural model, CaSE, to advance the state-of-the-art. We implement state-of-the-art methods from related tasks and conduct extensive experiments to show that: (1) the proposed CaSE model can achieve state-of-the-art performance; (2) the proposed STI and PPG modules can bring large improvements; and (3) SaaC is a more challenging dataset than previously introduced ones, leaving significant room for further improvements.

A limitation of this work is that the SaaC dataset is built upon the TREC CAsT dataset which adopts a “user asks - system responds” paradigm, and, hence, does not consider any mixed initiative aspects. It is also worth mentioning that we assume that each conversation session is well separated

in this work. However, in a practical system, a component is needed to detect when the users start a new session.

As to future work, on the one hand, we plan to further improve the performance of CaSE by proposing transfer learning methods in order to leverage more multi-turn conversational datasets (e.g., conversational QA [41, 63], conversational MRC [10] and chitchat [30]). On the other hand, we plan to study how to address the cases in SaaC when there is no correct answer or only a partially correct answer. Besides, it remains to be studied what progress we can make when different aspects are considered simultaneously in one system or model, e.g., asking clarifying questions [2, 67], conversational search and recommendation [11, 14, 50, 58], chitchats [61], conversation management [54] and so on.

CODE AND DATA

The SaaC dataset and the code of all the methods used for comparison in this paper are shared at <https://github.com/PengjieRen/CaSE-1.0>.

ACKNOWLEDGMENTS

We thank our anonymous reviewers and the guest editors for helpful feedback. This work was partially supported by the Netherlands Organisation for Scientific Research (NWO) under project nr CI-14-25, the NWO Innovational Research Incentives Scheme Vidi (016.Vidi.189.039), the NWO Smart Culture - Big Data / Digital Humanities (314-99-301), the H2020-EU.3.4. - SOCIETAL CHALLENGES - Smart, Green And Integrated Transport (814961), the Google Faculty Research Awards program, the Natural Science Foundation of China (61972234, 61902219, 61672324, 61672322, 62072279), the Key Scientific and Technological Innovation Program of Shandong Province (2019JZZY010129), the Tencent AI Lab Rhino-Bird Focused Research Program (JR201932), the Fundamental Research Funds of Shandong University, and the Innovation Center for AI (ICAI).

All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Mohammad Aliannejadi, Manajit Chakraborty, Esteban Andrés Rissola, and Fabio Crestani. 2020. Harnessing Evolution of Multi-Turn Conversations for Effective Answer Retrieval. In *Proceedings of 2020 Conference on Human Information Interaction and Retrieval (CHIIR'20)*.
- [2] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking Clarifying Questions in Open-Domain Information-Seeking Conversations. In *Proceedings of the 42Nd International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR'19)*. 475–484.
- [3] Leif Azzopardi, Mateusz Dubiel, Martin Halvey, and Jeffery Dalton. 2018. Conceptualizing Agent-human Interactions during the Conversational Search Process. In *The 2nd Workshop on Conversational Approaches to Information Retrieval*.
- [4] Nicholas J Belkin. 1980. Anomalous States of Knowledge as A Basis for Information Retrieval. *Canadian Journal of Information Science* 5, 1 (1980), 133–143.
- [5] Nicholas J Belkin, Colleen Cool, Adelheit Stein, and Ulrich Thiel. 1995. Cases, Scripts, and Information-seeking Strategies: On the Design of Interactive Information Retrieval Systems. *Expert Systems with Applications* 3 (1995), 379–395.
- [6] Niels O Bernsen, Hans Dybkjær, and Laila Dybkjær. 2012. *Designing Interactive Speech Systems: From First Ideas to User Testing*. Science & Business Media.
- [7] Keping Bi, Qingyao Ai, Yongfeng Zhang, and W. Bruce Croft. 2019. Conversational Product Search Based on Negative Feedback. In *Proceedings of the 28th ACM International Conference on Information & Knowledge Management (CIKM'19)*. 359–368.
- [8] Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - A Large-Scale Multi-Domain Wizard-of-Oz Dataset for Task-Oriented Dialogue Modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP'18)*. 5016–5026.

- [9] Ben Carterette, Virgil Pavlu, Hui Fang, and Evangelos Kanoulas. 2009. Million Query Track 2009 Overview. In *TREC 2009*.
- [10] 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 (EMNLP'18)*. Association for Computational Linguistics, 2174–2184.
- [11] Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards Conversational Recommender Systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'16)*. 815–824.
- [12] 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 (CIKM'19)*. 729–738.
- [13] W. Bruce Croft and R. H. Thompson. 1987. I3R: A New Approach to The Design of Document Retrieval Systems. *Journal of the Association for Information Science and Technology* 38, 6 (1987), 389–404.
- [14] Jeffrey Dalton, Chenyan Xiong, Vaibhav Kumar, and Jamie Callan. 2020. CASt-19: A Dataset for Conversational Information Seeking. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'20)*. 1985–1988.
- [15] Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, and Andrew McCallum. 2019. Multi-step Retriever-Reader Interaction for Scalable Open-domain Question Answering. In *International Conference on Learning Representations (ICLR'19)*.
- [16] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered Conversational Agents. In *The International Conference on Learning Representations (ICLR'19)*.
- [17] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL'18)*. 2–7.
- [18] 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 *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*. 2942–2951.
- [19] Alexandru Balan Gustavo Penha and Claudia Hauff. 2019. Introducing MANTIS: a novel Multi-Domain Information Seeking Dialogues Dataset. [arXiv:cs.CL/1912.04639](https://arxiv.org/abs/1912.04639)
- [20] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching Machines to Read and Comprehend. In *Proceedings of the 29th International Conference on Neural Information Processing Systems (NIPS'15)*. 1693–1701.
- [21] Ryuichiro Higashinaka, Kenji Imamura, Toyomi Meguro, Chiaki Miyazaki, Nozomi Kobayashi, Hiroaki Sugiyama, Toru Hirano, Toshiro Makino, and Yoshihiro Matsuo. 2014. Towards An Open-domain Conversational System Fully Based on Natural Language Processing. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING'14)*. 928–939.
- [22] Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL'17)*. 1601–1611.
- [23] Evangelos Kanoulas, Carterette Ben, Mark Hall, Paul Clough, and Mark Sanderson. 2011. Overview of the TREC 2011 Session Track. In *TREC 2011*.
- [24] Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. The NarrativeQA Reading Comprehension Challenge. *Transactions of the Association for Computational Linguistics* 6 (2018), 317–328.
- [25] Bernhard Kratzwald and Stefan Feuerriegel. 2019. Learning from On-Line User Feedback in Neural Question Answering on the Web. In *The World Wide Web Conference (WWW'19)*. 906–916.
- [26] Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent Retrieval for Weakly Supervised Open Domain Question Answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL'19)*. 6086–6096.
- [27] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-action-reflection: Towards deep interaction between conversational and recommender systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM'20)*. 304–312.
- [28] Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL'18)*. 1437–1447.
- [29] James Lester, Karl Branting, and Bradford Mott. 2004. Conversational Agents. *The Practical Handbook of Internet Computing* (2004), 220–240.

- [30] Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The Ubuntu Dialogue Corpus: A Large Dataset for Research in Unstructured Multi-Turn Dialogue Systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '15)*. 285–294.
- [31] Chuan Meng, Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. RefNet: A Reference-aware Network for Background based Conversation. In *The 34th AAAI Conference on Artificial Intelligence (AAAI'20)*.
- [32] Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M Khapra. 2018. Towards Exploiting Background Knowledge for Building Conversation Systems. In *The 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP'18)*. 2322–2332.
- [33] 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*.
- [34] Kyosuke Nishida, Itsumi Saito, Kosuke Nishida, Kazutoshi Shinoda, Atsushi Otsuka, Hisako Asano, and Junji Tomita. 2019. Multi-style Generative Reading Comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL'19)*. 2273–2284.
- [35] Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A Deep Reinforced Model for Abstractive Summarization. In *International Conference on Learning Representations (ICLR'18)*.
- [36] 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 (SIGIR'20)*. 539–548.
- [37] Chen Qu, Liu Yang, W. Bruce Croft, Johanne R. Trippas, Yongfeng Zhang, and Minghui Qiu. 2018. Analyzing and Characterizing User Intent in Information-seeking Conversations. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR'18)*. ACM, 989–992.
- [38] Chen Qu, Liu Yang, W. Bruce Croft, Yongfeng Zhang, Johanne R. Trippas, and Minghui Qiu. 2019. User Intent Prediction in Information-Seeking Conversations (*CHIIR'19*). 25–33.
- [39] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval (CHIIR'17)*. 117–126.
- [40] 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 (EMNLP'16)*. 2383–2392.
- [41] 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.
- [42] Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-local Knowledge Selection for Background based Conversation. In *The 34th AAAI Conference on Artificial Intelligence (AAAI'20)*.
- [43] Pengjie Ren, Zhumin Chen, Zhaochun Ren, Furu Wei, Jun Ma, and Maarten de Rijke. 2017. Leveraging Contextual Sentence Relations for Extractive Summarization Using a Neural Attention Model. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'17)*. 95–104.
- [44] Amrita Saha, Vardaan Pahuja, Mitesh M. Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018. Complex Sequential Question Answering: Towards Learning to Converse Over Linked Question Answer Pairs with a Knowledge Graph. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. 705–713.
- [45] Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL'17)*. 1073–1083.
- [46] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hananneh Hajishirzi. 2017. Bidirectional Attention Flow for Machine Comprehension. In *International Conference on Learning Representations (ICLR'17)*.
- [47] Igor Shalyminov, Alessandro Sordani, Adam Atkinson, and Hannes Schulz. 2020. Fast Domain Adaptation for Goal-Oriented Dialogue Using a Hybrid Generative-Retrieval Transformer. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing*. 8039–8043.
- [48] Igor Shalyminov, Alessandro Sordani, Adam Atkinson, and Hannes Schulz. 2020. Hybrid Generative-Retrieval Transformers for Dialogue Domain Adaptation. *CoRR* abs/2003.01680 (2020).
- [49] Tao Shen, Xiubo Geng, Tao Qin, Daya Guo, Duyu Tang, Nan Duan, Guodong Long, and Daxin Jiang. 2019. Multi-Task Learning for Conversational Question Answering over a Large-Scale Knowledge Base. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP'19)*. 2442–2451.
- [50] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'18)*. 235–244.

- [51] Chuanqi Tan, Furu Wei, Nan Yang, Bowen Du, Weifeng Lv, and Ming Zhou. 2018. S-Net: From Answer Extraction to Answer Synthesis for Machine Reading Comprehension. In *The 32nd AAAI Conference on Artificial Intelligence (AAAI'18)*. 5940–5947.
- [52] Paul Thomas, Mary Czerwinski, Daniel McDuff, Nick Craswell, and Gloria Mark. 2018. Style and Alignment in Information-Seeking Conversation. In *Proceedings of the 2018 Conference on Human Information Interaction and Retrieval (CHIIR'18)*. 42–51.
- [53] Johanne R. Trippas, Damiano Spina, Lawrence Cavedon, Hideo Joho, and Mark Sanderson. 2018. Informing the Design of Spoken Conversational Search: Perspective Paper. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval (CHIIR'18)*. 32–41.
- [54] Johanne R. Trippas, Damiano Spina, Paul Thomas, Mark Sanderson, Hideo Joho, and Lawrence Cavedon. 2020. Towards a model for spoken conversational search. *Information Processing & Management* 57, 2 (2020).
- [55] Svitlana Vakulenko, Ilya Markov, and Maarten de Rijke. 2017. Conversational Exploratory Search via Interactive Storytelling. In *1st International Workshop on Search-Oriented Conversational AI*.
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. 6000–6010.
- [57] Nikos Voskarides, Dan Li, Andreas Panteli, and Pengjie Ren. 2019. ILPS at TREC 2019 Conversational Assistant Track. In *TREC 2019*.
- [58] 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 (SIGIR'20)*. 921–930.
- [59] Alexandra Vtyurina, Denis Savenkov, Eugene Agichtein, and Charles L. A. Clarke. 2017. Exploring Conversational Search With Humans, Assistants, and Wizards. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA'17)*. 2187–2193.
- [60] Xiangnan He, Yisong Miao, Xiang Wang, Liang Chen, Wenqiang Lei, Gangyi Zhang, and Tat-Seng Chua. 2020. Interactive Path Reasoning on Graph for Conversational Recommendation. In *Proceedings of the 2020 ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'20)*.
- [61] Wei Wu and Rui Yan. 2018. Deep Chit-Chat: Deep Learning for ChatBots. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*.
- [62] Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic Coattention Networks For Question Answering. In *International Conference on Learning Representations (ICLR'17)*.
- [63] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP'18)*. 2369–2380.
- [64] Wen-tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. The Value of Semantic Parse Labeling for Knowledge Base Question Answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL'16)*. 201–206.
- [65] Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. 2016. Neural Generative Question Answering. In *Proceedings of the Workshop on Human-Computer Question Answering*.
- [66] Hamed Zamani and Nick Craswell. 2019. Macaw: An Extensible Conversational Information Seeking Platform. *arXiv preprint arXiv:1912.08904* (2019).
- [67] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating Clarifying Questions for Information Retrieval. In *Proceedings of the 29th International Conference on World Wide Web (WWW'20)*.
- [68] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards Conversational Search and Recommendation: System Ask, User Respond. In *Proceedings of the 27th ACM International Conference on Information & Knowledge Management (CIKM'18)*. 177–186.