

# Wizard of Search Engine

Access to Information through Conversations with Search Engines

Pengjie Ren<sup>1</sup> Zhongkun Liu<sup>1</sup> Xiaomeng Song<sup>1</sup> Hongtao Tian<sup>1</sup>

Zhumin Chen<sup>1\*</sup> Zhaochun Ren<sup>1\*</sup> Maarten de Rijke<sup>2,3</sup>

<sup>1</sup>Shandong University, Qingdao, China

<sup>2</sup>University of Amsterdam, Amsterdam, The Netherlands

<sup>3</sup>Ahold Delhaize, Zaandam, The Netherlands

{renpengjie,chenzhumin,zhaochun.ren}@sdu.edu.cn,{m.derijke}@uva.nl

## ABSTRACT

Conversational information seeking (CIS) is playing an increasingly important role in connecting people to information. Due to a lack of suitable resources, previous studies on CIS are limited to the study of conceptual frameworks, laboratory-based user studies, or a particular aspect of CIS (e.g., asking clarifying questions).

In this work, we make three main contributions to facilitate research into CIS: (1) We formulate a pipeline for CIS with six subtasks: intent detection, keyphrase extraction, action prediction, query selection, passage selection, and response generation. (2) We release a benchmark dataset, called *wizard of search engine* (WISE), which allows for comprehensive and in-depth research on all aspects of CIS. (3) We design a neural architecture capable of training and evaluating both jointly and separately on the six sub-tasks, and devise a pre-train/fine-tune learning scheme, that can reduce the requirements of WISE in scale by making full use of available data.

We report useful characteristics of the CIS task based on statistics of the WISE dataset. We also show that our best performing model variant is able to achieve effective CIS. We release the dataset, code as well as evaluation scripts to facilitate future research by measuring further improvements in this important research direction.

## CCS CONCEPTS

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

## KEYWORDS

Search engine; Conversational search; Dataset

### ACM Reference Format:

Pengjie Ren, Zhongkun Liu, Xiaomeng Song, Hongtao Tian, Zhumin Chen, Zhaochun Ren, and Maarten de Rijke. 2021. Wizard of Search Engine: Access to Information through Conversations with Search Engines. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development*

\*Corresponding author.

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 the author(s) 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](mailto:permissions@acm.org).

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3462897>

in *Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3404835.3462897>

## 1 INTRODUCTION

Search engines have become the dominant way for people around the world to interact with information. Today, traditional search engines (TSEs) have settled on a query-SERP (Search Engine Result Page) paradigm, where a searcher issues a query in the form of keywords to express their information need, and the search engine responds with a SERP containing a ranked list of snippets. Even though this has become the dominant paradigm for interactions with search engines, it is confined to clear navigational or transactional intents where the user wants to reach a particular site [7], or relatively simple informational intents, where the user's questions can be answered with short text spans [8]. However, it is not effective when the user has complex and composite intents, especially when their information need is unclear and/or exploratory [43].

*Conversational information seeking* (CIS) has emerged as a new paradigm for interactions with search engines [3, 11, 47]. In contrast to the query-SERP paradigm, CIS allows users to express their information need by directly conducting conversations with search engines. In this way, CIS systems can better capture a user's intent by taking advantage of the flexibility of mixed-initiative interactions, and provide useful information more directly by returning human-like responses. CIS is increasingly attracting attention. Currently, research on CIS is mostly focused on the following aspects: asking clarifying questions [2, 48], conversational search [13, 41], conversational question answering [29, 31], and conversational recommendation [9, 34]. A complete CIS system should be a mixture of these aspects and go far beyond these [16]. Other aspects, e.g., chitchat [45], user intent understanding [27], feedback analysis [19], conversation management, and so on [38], are still not well studied in the context of CIS scenario. Since no complete CIS models capable of handling all these aspects have been developed so far, there is a lack of comprehensive analysis on the performance of those aspects when achieved and/or evaluated simultaneously. This is largely because there is no practical formalism or suitable resource for CIS. As a result, related research is restricted to user studies [36, 42] and/or theoretical/conceptual CIS frameworks [4, 28].

We contribute to CIS by pursuing three goals in this work: a pipeline, a dataset, and a model. Let us expand on each of these.

(1) We formulate a pipeline for CIS consisting of six sub-tasks: intent detection (ID), keyphrase extraction (KE), action prediction (AP), query selection (QS), passage selection (PS), and response generation (RG), as shown in Figure 1. ID identifies the

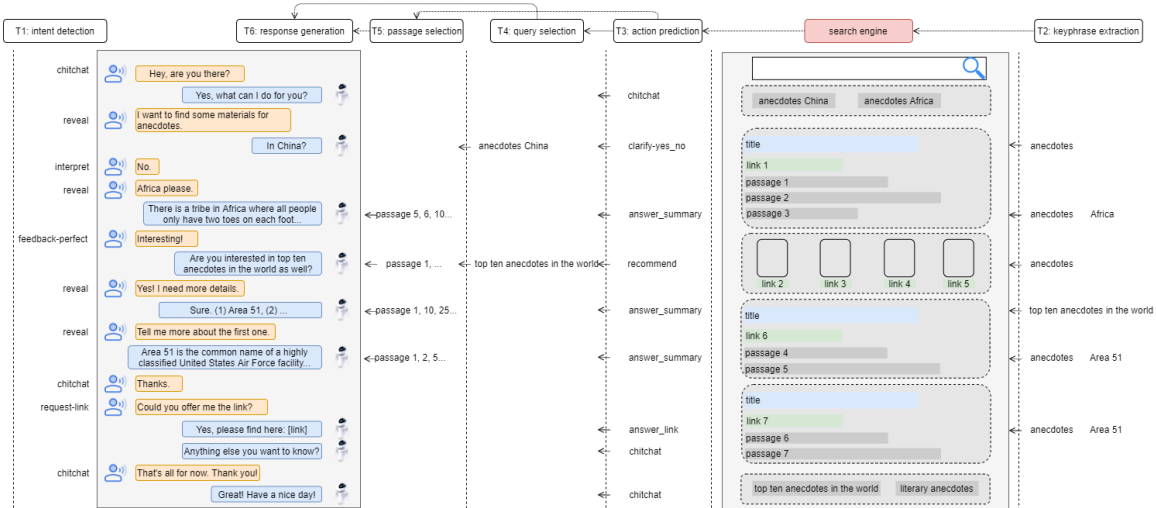


Figure 1: Pipeline formalism for conversational information seeking (CIS).

general intent behind the current user utterance, e.g., greetings (“chitchat”) or revealing information need (“reveal”). Then, KE extracts keyphrases from previous conversations to represent the fine-grained information need behind the current user utterance (if any), which can be used to retrieve candidate queries (i.e., suggested queries, related queries) and candidate documents/passages. AP is used for conversation management as it decides the system action at each turn. QS and PS select the relevant queries and passages from which RG can distill useful information (if any) to generate conversational responses. We make a close connection to TSEs so as to retain its characteristics as much as possible (e.g., candidate queries and passage retrieval) while introducing new characteristics of CIS (e.g., conversation management, conversational response generation).

- (2) We build a dataset of human-human CIS conversations in a wizard-of-oz fashion, called *wizard of search engine* (WISE), where two workers play the role of seeker and intermediary, respectively. We first collect search sessions from a commercial search engine, which cover a diverse range of information needs. Then, we write guidelines for the seekers by inferring the intents behind the keyword queries in the search sessions, e.g., “[anecdotes, anecdotes in the world, Tolstoy, Tolstoy Wikipedia] → You want to find some materials about anecdotes. Describe the features of anecdotes you are interested in. Discover the introduction of the anecdotes and ask for links.” The seekers can better understand the information need when provided with guidelines than keyword queries from search sessions. We make sure the guidelines cover diverse search scenarios, e.g., when the users have specific, unclear, or multiple intents. Finally, we build a conversational interface for the workers to conduct conversations. The interface for the intermediaries is backed by a commercial search engine. We record all the operations the workers take, which can be used for supervised learning or direct evaluation for all six sub-tasks.
- (3) We propose a modularized end-to-end neural architecture, where each module corresponds to one of the six sub-tasks listed above. Thus, we have a natural way to train and evaluate the model,

Table 1: Summary of main notation used in the paper.

$X^\tau$	User utterance at turn $\tau$ . $X_t^\tau$ represents the token at position $t$ .
$Y^\tau$	System utterance at turn $\tau$ . $Y_t^\tau$ represents the token at position $t$ .
$C$	A conversation. $C^{:\tau} = \dots, X^{\tau-1}, Y^{\tau-1}$ represents the utterances until turn $\tau - 1$ .
$d$	A passage $d \in \mathbb{D}$ . $D^\tau$ represents the set of candidate passages at turn $\tau$ .
$q$	A query $q \in \mathbb{Q}$ . $Q^\tau$ represents the set of candidate queries at turn $\tau$ .
$i$	A user intent $i \in \mathbb{I}$ .
$a$	A system action $a \in \mathbb{A}$ .
$k$	A keyphrase. $K^\tau$ is the set of keyphrases for $C^{:\tau} + X^\tau$ .
$h$	A hidden representation $h \in H$ . $H$ is a sequence.

both jointly and separately, by either forwarding the ground truth or the prediction to downstream modules. We also devise a pre-train/fine-tune learning scheme to make full use of available data and reduce the requirements of supervised learning. We use four related tasks/datasets to guarantee that all the modules can be pre-trained.

We conduct analyses to determine several useful characteristics of CIS as reflected in the WISE dataset. We also carry out experiments to analyze the performance of different model variants on the six sub-tasks, as well as the effectiveness of the pre-training tasks. We show that our best performing model variant is able to achieve effective CIS. To facilitate experimental reproducibility, we share the dataset, code, and evaluation scripts.

## 2 TASKS

Table 1 summarizes the notation used in the paper.

### 2.1 Task 1: Intent detection (ID)

The goal of this task is to identify the general intent of each user utterance. There have been several similar tasks [27]. However, the

proposed user intent taxonomies are either too general or not well mapped to user behavior in TSEs. Considering previous work [4, 28] while making a connection to TSEs, we define a new user intent taxonomy, as shown in Table 2. Most intents can map to the corresponding user behavior or operation in TSEs, except for “chitchat” and “request-rephrase”, because TSEs do not have conversations and the outputs are always organized as SERPs. The algorithmic problem of ID in CIS is defined as learning a mapping:  $\{C^\tau, X^\tau\} \rightarrow i$ . That is, given the context  $C^\tau$  and the current user utterance  $X^\tau$ , the goal is to predict the user intent  $i$  of for the utterance  $X^\tau$ .

## 2.2 Task 2: Keyphrase extraction (KE)

Conversations usually contain more content than necessary in order to ensure the integrity of grammar and semantics, etc. Ellipsis and anaphora are common in conversations, so the current user utterance itself is usually not semantically complete. This is especially true as conversations go on and when the users are not familiar with the topic, or unclear about the information goals themselves. It is challenging to fetch documents that are relevant to the intent in the current user utterance by simply issuing the current user utterance to the TSE. To this end, similar to [41], we define KE as a task to extract keyphrases from the current user utterance as well as the previous conversational context, which can help TSEs to retrieve relevant documents for the current user utterance. Formally, KE is defined as learning a mapping:  $\{C^\tau, X^\tau\} \rightarrow K^\tau$ .

## 2.3 Task 3: Action prediction (AP)

To connect users to the right information with maximal utility, a good CIS system should be able to take appropriate actions at the right time, e.g., helping the users to clarify their intents whenever unclear, and providing information in the right form by tracking user expectations. AP has been well studied in task-oriented dialogue systems (TDSs) in a closed-domain setting (a.k.a. dialogue policy) [26]. However, AP is not yet well addressed in CIS. Although there have been some investigations about what system actions are necessary in CIS, to the best of our knowledge, it has not yet been put into practice [4, 28]. By recalling the operations in TSEs, we introduce a system action taxonomy in Table 3. The AP task is formulated as learning a mapping:  $\{C^\tau, X^\tau, Q^\tau, D^\tau\} \rightarrow a$ . Note that this task also takes the candidate queries  $Q^\tau$  and candidate passages  $D^\tau$  fetched from TSEs as input, as they will influence the action the CIS system should take. For example, to decide whether “no-answer” is the action to be taken, the CIS system must check whether the answer is available in  $Q^\tau$  and/or  $D^\tau$ .

## 2.4 Task 4: Query selection (QS)

In TSEs, a list of suggested or related queries is usually presented to users to help them clarify intents or find similar intents. In CIS, we cannot always simply list the suggested or related queries to users, as this is not natural in conversations. Previous studies cast this problem as a problem of asking clarifying questions, and model it as either selecting a question from a question pool prepared in advance [2] or generating a clarifying question for a given keyword query [48]. However, neither option fits in our formulation. Besides, how to recommend similar search intents rather than asking clarifying questions is neglected. We introduce the QS task as a

step in the CIS pipeline. Depending on the predicted actions from the upstream AP task, QS aims to select queries fetched from TSEs that can be used by the downstream task to generate clarifying questions or recommend similar search intents. Formally, QS is defined as learning a mapping:  $\{C^\tau, X^\tau, Q^\tau, (i, a)\} \rightarrow Q_s^\tau$ , where  $Q^\tau$  is set of suggested or related queries fetched from TSEs with the keyphrases from KE as queries, and  $i$  and  $a$  are the user intent and system action from ID and AP, respectively.

## 2.5 Task 5: Passage selection (PS)

Passage or document ranking has long been studied in TSEs [46]. Recently, this has also been investigated in a conversational context [13]. The situation is a bit different for PS in CIS where the context is not only about the proactive questions from users but also the clarifying questions or recommendations from the systems. And the goal is not always to find the relevant passages for the user questions but also non-relevant passages that contain related topics in order to do clarification or recommendation sometimes [43]. Given the same input as in QS, the PS task is to select passages based on which to generate system responses by the downstream task:  $\{C^\tau, X^\tau, D^\tau, (i, a)\} \rightarrow D_s^\tau$ , where  $D^\tau$  is a set of passages fetched from a TSE with the same keyphrases from KE as queries.

## 2.6 Task 6: Response generation (RG)

RG is a fundamental component for many NLP tasks, e.g., chitchat [30], TDS [25], conversational question answering [6]. The goals or formulations of RG vary across tasks. The role of RG in our formulation of CIS is to translate system actions and predictions of the above tasks into natural language responses. For example, if the system action is “clarify”, the RG part needs to generate a clarifying question by referring to the selected queries. In this work, we only focus on free text responses. The formulation of RG is defined as a learning problem:  $\{C^\tau, X^\tau, Q^\tau, D^\tau, (i, a, Q_s^\tau, D_s^\tau)\} \rightarrow Y^\tau$ .

## 3 DATASET

The WISE dataset is built based on the following setting: two participants engage in a CIS conversation, one plays the role of a searcher while the other plays the role of a knowledgeable expert (referred to as wizard). At each stage of the conversation, the searcher talks to the wizard freely. Their goal is to follow the general instruction about a chosen search intent, go into depth about it, and try to get all information that has been mentioned in the instruction. The wizard helps the searcher to find the information by interacting with a search engine. Their goal is to help the searcher (1) clarify their search intents whenever unclear, (2) answer the searcher’s questions by finding and summarizing the relevant information in search results, and (3) recommend information that has not been asked but is related to the searcher’s search intents or other related topics the searcher might be interested in. There are two steps in building the WISE dataset, as shown in Figure 2: collecting search intents, and collecting CIS conversations.

### 3.1 Collecting search intents

First, we collect a set of 1,196 search intents from a search log of a commercial search engine. Each search intent is based on a specific search session. For a given search session, e.g., “[anecdotes,

**Table 2: User intent for the seeker.**

Intent	Explanation	Example	TSE operations
reveal	Reveal a new intent, or refine an old intent proactively.	User: I want to see a movie. (reveal) User: Can you tell me more about it? (reveal)	Issue a new query.
revise	Revise an intent proactively when there is wrong expression, e.g., grammatical issues, unclear expression.	User: Tell me some non-dairy milks. User: I mean dairy not diary. (revise)	Revise the query.
interpret	Interpret or refine an intent by answering a clarification question from the system.	User: Do you know The Avengers? System: Do you mean the movie, novel or game? User: The movie (interpret)	Select suggested queries.
request-rephrase	Request the system to rephrase the response if it is not understandable.	Sorry, I didn't get it. (request-rephrase)	-
chitchat	Greetings or other utterances that are not related to the information need.	I see. (chitchat) Are you there? (chitchat)	-

**Table 3: System actions for the intermediary.**

Action	Explanation	Example	TSE operations
clarify	yes-no	Do you want to the plot? (clarify-yes-no)	Suggest queries.
	choice	Ask questions to clarify user intent when it is unclear or exploratory. Do you want to know its plot, cast or director? (clarify-choice)	
	open	What information do you want to know? (clarify-open)	
answer-type	opinion	Give advice, ideas, suggestions, or instructions. The response is more subjective. I recommend xxx, because ... (answer-opinion)	Provide results.
	fact	Give a single, unambiguous answer. The response is objective and certain. Her birthday is xxx. (answer-fact)	
	open	Give an answer to an open-ended question, or one with unconstrained depth. The response is objective but may be different depending on the perspectives. One of the reasons of the earthquake is that... (answer-open)	
answer-form	free-text	The disadvantages of Laminate Flooring are that ..... (answer_free_text)	
	list	Answer the user intent by providing information in the right form or when being asked to answer in a particular form. Area 51. ... (answer_list)	
	steps	1. Click on ... 2. (answer_steps)	
	link	You can find the video here: [link]. (answer_link)	
no-answer	If there is no relevant information found, notice the user.	Sorry, I cannot find any relevant information. (no-answer)	No answer found.
request-rephrase	Ask the user to rephrase its question if it is unclear.	I didn't really get what you mean. (request-rephrase)	-
chitchat	Greetings or other content that are not related to the information need.	Hi. (chitchat) Yes, I am ready to answer your questions. (chitchat)	-

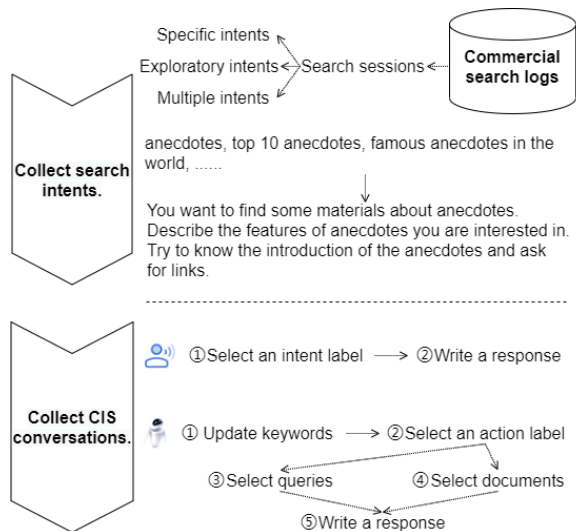


Figure 2: Pipeline for data construction.

anecdotes in the world, Tolstoy, Tolstoy Wikipedia, Tolstoy movie]”, the workers are asked to infer the search intent behind it using their imagination, and write down a description in one of the following forms: (1) specific intent, e.g., “You are interested in anecdotes about famous writers in the world like Tolstoy. Try to learn about his/her basic information, life story and representative works, and so on.” (2) exploratory intent, e.g., “You want to find some material about anecdotes. Describe the features of anecdotes you are interested in. Try to learn about the introduction of the anecdotes and ask for links.” (3) multiple intents, e.g., “(1) or (2) + You also want to find a movie that is based on his/her works. Ask for information about the director, plot, links and so on.”

### 3.2 Collecting CIS conversations

The conversation flows as follows. (1) The searcher randomly picks a search intent and starts by asking questions directly or by uttering a greeting. When sending a message, the searcher is required to select a label from Table 2 that can best describe the intent of the message. (2) Whenever the wizard receives the message from the searcher, s/he needs to extract keyphrases from the conversation history that are used to fetch results from a search engine. Then s/he has to select a label from Table 3 that reflects the action s/he is going to take. After that, s/he needs to select the relevant queries and/or documents, based on which to formulate the responses. (3) The conversation repeats until the searcher ends the chat (after a minimum of 7 turns each). For each turn, both the searcher and wizard can send multiple messages at once by clicking the “send another message” option. After collecting data in this way, the goal is to then replace the wizard with a learned CIS system that will speak to a human searcher instead.

### 3.3 Measures/Disclaimers for data quality

3.3.1 *Measures.* We took several measures to guarantee the data quality. First, we provided detailed documentation on the task, the definition of the labels, the steps involved in creating the data, and the instruction of using the interfaces. Second, we recorded a video

Table 4: Datasets statistics.

	train	valid	test	
			test(seen)	test(unseen)
conversations	705	200	500	500
turns	6,099	1,925	4,860	4,626
utterances	13,569	4,125	10,410	9,852
avg turns	8.65	9.63	9.72	9.25
avg utterances	19.25	20.63	20.82	19.70
avg words	27.06	29.3	25.94	28.3

to demonstrate the whole process and enhance the considerations to take into account. Third, we also maintained a list of positive and negative examples from us as well as the workers as a reference for the other workers. Fourth, after obtaining the data, we manually went over it and corrected the following issues: grammatical issues, mis-labels, long responses without summarization.

3.3.2 *Disclaimers.* In all the provided instruction materials, we described the purpose of this data construction effort and pointed out that the data will only be used for research. We did not record any information about the workers and warned the workers not to divulge any of their private information during conversations. We filtered out all adult content and/or offensive content that might raise ethical issues when collecting the search intents, and asked the workers to be careful about such content in conversations too.

### 3.4 Statistics

It took 24 workers 3 months to build the dataset. The final dataset we collected consists of 1,905 conversations with 37,956 utterances spread over 1,196 search intents. In total, the dataset contains 12 different Intents and 23 different Actions, covering a variety of conversation topics such as literature, entertainment, art, etc. Each conversation contains 7 to 42 utterances. The average number of utterances per conversation is 19.9, and the average number of turns per conversation is 9.2. Each utterance has an average of 27.3 words. We divide the data into 705 conversations for training, 200 conversations for validation, and 1,000 for testing. The test set is split into two subsets, test(seen) and test(unseen). The test(seen) dataset contains 442 overlapping search intents with the training set with new conversations. Test(unseen) consists of 500 search intents never seen before in train or validation. The overall data statistics can be found in Table 4.

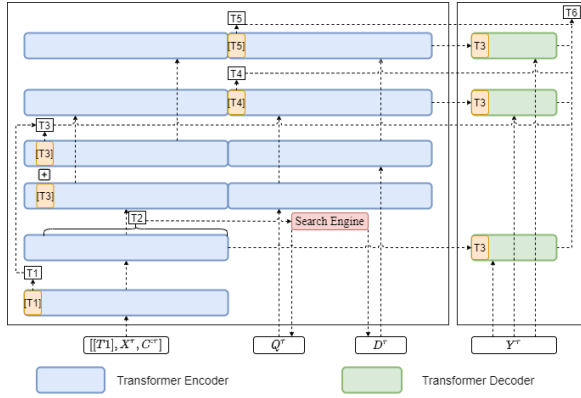
## 4 METHODOLOGY

We propose a model to replace the wizard in data construction and devise a pre-train/fine-tune learning scheme to train the model.

### 4.1 Model

As shown in Figure 3, we propose an architecture that relies mainly on transformer encoders and decoders.

4.1.1 *Transformer encoder.* The transformer encoder consists of two layers: a self-attention layer and a position-wise feed-forward layer. The self-attention layer adopts the multi-head attention mechanism [39] to identify useful information in the input sequence. The feed-forward layer consists of two linear transformations with



**Figure 3: Model architecture.**

a ReLU activation in between [1]. Before each layer, layer normalization is applied [5]. After each layer, a residual connection is applied. For any given sequence  $S$ , the transformer encoder outputs a sequence of hidden representations  $H \in \mathbb{R}^{l_S \times dim}$  corresponding to  $S$ :  $H = \text{TEncoder}(S)$ , where  $l_S$  is the sequence length,  $dim$  is the hidden size, and  $\text{TEncoder}$  is short for transformer encoder.

**4.1.2 Transformer decoder.** The transformer decoder consists of three layers: an output-attention layer, an input-attention layer and a position-wise feed-forward layer. They all use the same implementation as in the transformer encoder. The difference is that the output-attention and input-attention layers try to identify useful information in the output sequence and input sequence, respectively. For the output sequence, a look-ahead mask is used to mask the future positions so as to prevent information from future positions from affecting current positions [39]. Again, before and after each layer, a layer normalization and a residual connection are applied, respectively. For any given output and input sequences  $S_O$  and  $S_I$ , the transformer decoder outputs a sequence of hidden representations  $H \in \mathbb{R}^{l_O \times dim}$  corresponding to  $S_O$  and the input-attention weights  $W \in \mathbb{R}^{l_O \times l_I}$  corresponding to  $S_I$ :  $H, W = \text{TDecoder}(S_O, S_I)$ , where  $l_O$  and  $l_I$  are the output and input sequence lengths, respectively.  $\text{TDecoder}$  is short for transformer decoder.

**4.1.3 Task 1: Intent detection (ID).** We concatenate  $C^\tau$  and  $X^\tau$  in reverse order of utterance turns, and put a special token “[T1]” at the beginning to get the input of Task 1:  $S_{T1} = [[T1], X^\tau, C^\tau]$ . The sequence goes to a token and positional embedding layer to get  $E_C$ . Then, we input the embedding sequence into a stack of transformer encoders to get hidden representations for all tokens:  $H_{T1} = \text{TEncoder}(E_C)$ . Finally, we get the hidden representation w.r.t. “[T1]” and use a linear classifier with softmax to do ID.

**4.1.4 Task 2: Keyphrase extraction (KE).** We model KE as a binary classification on each token of  $C^\tau$  and  $X^\tau$ . To do so, the same stack of transformer encoders as for ID is used to get the hidden representations for KE based on the representations from ID:  $H_{T2} = \text{TEncoder}(H_{T1})$ . A linear classifier with sigmoid is used to predict whether each token belongs to a keyphrase or not.

**4.1.5 Task 3: Action prediction (AP).** To predict the system action, we also need to consider the query and passage candidates  $Q^\tau$  and  $D^\tau$ . For each  $q$  or  $d$ , we first input it into a token and

positional embedding layer to get  $E_q$  or  $E_d$ . We then concatenate it with the hidden representations from KE and fuse the context and query/passage information with the same stack of transformer encoders as for ID:  $H_{T3}^q = \text{TEncoder}([H_{T2}, E_q])$  and  $H_{T3}^d = \text{TEncoder}([H_{T2}, E_d])$ . Note that another special token “[T3]” is used (see Figure 3). A hidden representation corresponding to “[T3]” is obtained for each  $q$  or  $d$ . We combine them with a max pooling to get a single hidden representation, on top of which a linear classifier with softmax is applied to do AP.

**4.1.6 Task 4: Query selection (QS).** For QS, a similar architecture as AP is adopted. The same stack of transformer encoders is used to get hidden representations for QS based on the hidden representations  $H_{T3}^q$  from AP:  $H_{T4} = \text{TEncoder}(H_{T3}^q)$ . We get the hidden representation corresponding to “[T4]”, on top of which a binary classifier with sigmoid is used to predict whether a query  $q$  should be selected or not.

**4.1.7 Task 5: Passage selection (PS).** Similar to QS, PS uses the same stack of transformer encoders to get hidden representations based on the hidden representations  $H_{T3}^d$  from AP:  $H_{T5} = \text{TEncoder}(H_{T3}^d)$ . The hidden representation corresponding to “[T5]” is used to predict whether a passage  $d$  should be selected or not using a binary classifier with sigmoid.

**4.1.8 Task 6: Response generation (RG).** We use a stack of transformer decoders to do RG. The output sequence in the transformer decoder is the ground truth response  $Y^\tau$  during training or the generated response  $\hat{Y}^\tau$  during inference. Note that we put the system action at the beginning to indicate the type of the response. We get the fused hidden representations by taking the context, query, and passage hidden representations from KE, QS, and PS as the input sequence in the transformer decoder, respectively. Specifically,  $H_{T6}^C, W_{T6}^C = \text{TDecoder}(E_Y, H_{T2})$ ;  $H_{T6}^q, W_{T6}^q = \text{TDecoder}(E_Y, H_{T4})$ ;  $H_{T6}^d, W_{T6}^d = \text{TDecoder}(E_Y, H_{T5})$ .  $H_{T6}^C, H_{T6}^q$  and  $H_{T6}^d$  are fused again with max pooling. Then a linear classifier with softmax is applied on top to predict the probability of the token at each time step  $P_g(Y_t^\tau)$ . We also extend the pointer mechanism [40] to generate a token by copying a token in the context, query or passage [33]. Pointer mechanisms are commonly used in natural language generation tasks. The copying probability is modeled as  $P_c(Y_t^\tau) = \max(W_{T6}^C(Y_t^\tau), P(q)W_{T6}^q(Y_t^\tau), P(d)W_{T6}^d(Y_t^\tau))$ , where  $W_{T6}^C(Y_t^\tau)$  is the sum of attention weights for  $Y_t^\tau$  in the context. A similar definition applies to  $W_{T6}^q(Y_t^\tau)$  and  $W_{T6}^d(Y_t^\tau)$ .  $P(q)$  and  $P(d)$  are the selection probability from QS and PS, respectively. The final probability  $P(Y_t^\tau)$  is a linear combination of  $P_g(Y_t^\tau)$  and  $P_c(Y_t^\tau)$ .

## 4.2 Learning

We use binary cross-entropy loss for KE, QS and PS and cross-entropy loss for ID, AP and RG. To solve complex tasks with a deep model, a large volume of data is needed. We get around this problem by first pre-training the model on four datasets from related tasks, as shown in Table 5. WebQA is a dataset for single-turn question answering. For each sample, it provides a simple question in natural language, a set of candidate passages, and an entity answer. DuReader is for single-turn machine reading comprehension. The questions are mostly complex and the answers are summaries of

Table 5: Pre-training datasets.

Datasets	T1 (ID)	T2 (KE)	T3 (AP)	T4 (QS)	T5 (PS)	T6 (RG)
WebQA [20]	✗	✗	✗	✗	✓	✓
DuReader [17]	✓	✗	✓	✗	✓	✓
DuConv [44]	✗	✓	✗	✓	✗	✓
KdConv [49]	✗	✓	✗	✓	✗	✓

the relevant passages. It also has question and answer types, which can be used for the pre-training of T1 and T3. DuConv and KdConv are for knowledge grounded conversations. DuConv has structured knowledge in triples while KdConv has free text sentences. Since the sentences in KdConv are usually short, we use it to pre-train T4. DuConv also has annotations of topic words in conversations which can be used to pre-train T2. For KdConv, we pre-train T2 by predicting the overlapping tokens in conversational contexts and answers. We do not pre-train T2 using WebQA or DuReader, because they are single-turn and the questions are very short usually. After pre-training, we fine-tune the model on the WISE data.

## 5 EXPERIMENTAL SETUP

*Evaluation metrics.* We adopt different evaluation metrics for different subtasks. We use BLEU-1 [24] and ROUGE-L [21] for keyphrase extraction (KE) and response generation (RG), which are commonly used in NLG tasks, e.g., QA, MRC [18, 23]. We report macro Precision, Recall and F1 for intent detection (ID), action prediction (AP), query selection (QS) and passage selection (PS).

*Implementation details.* For a fair comparison, we implement all models used in our experiments based on the same code framework. We set the word embedding size and hidden size to 512. We use the same vocabulary for all methods and its size is 21,129. The learning rate was increased linearly from zero to  $2.5 \times 10^{-4}$  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  $\epsilon = 10^{-8}$ ). We use 4 transformer layers for the encoder and 2 for the decoder, where the number of attention heads is 8. In WISE, the parameters of the transformer encoder, transformer decoder, and word embedding for each component of our model are shared. We train all models for 50 epochs, evaluate on a validation set for the best model in terms of BLEU-1 in RG task.

## 6 RESULT AND ANALYSIS

### 6.1 Performance analysis

We report the results of all methods on WISE; see Table 6. We consider three settings: (1) With pretraining (i.e., WISE in Table 6). (2) Without pretraining (i.e., WISE-pretrain in Table 6). (3) With pretraining and using the ground truth of ID, KE, AP, QS and PS (i.e., WISE+GT in Table 6). We also report the performance of WISE on the test (unseen) and test (seen) datasets, and on different actions. The results are shown in Table 7 and 8.

From Table 6, there are two main observations. First, WISE+GT achieves the best performance in terms of all key metrics. As we can see in Table 6, WISE+GT achieves 20.4% and 24.2% in terms of BLEU and ROUGE for RG, outperforming the strongest model

WISE by 5.8% and 7.3%, respectively. This shows that the human labels of ID, KE, AP, QS, PS do benefit RG.

Second, WISE achieves a higher performance than WISE-pretrain in terms of most metrics, indicating that the pretrain phrase is of great benefit for the performance of WISE. For example, the BLEU and ROUGE score of WISE for RG are 13.6% and 15.9%, much higher than those of WISE-pretrain, respectively. The reason is that the model can leverage external knowledge in the pretrain phrase, which boosts the final performance.

From Table 8, we obtain three main findings. First, for the action ‘Clarify’, the best performance in terms of KE, AP and QS is achieved w.r.t. ‘Clarify choice’ and in terms of RG w.r.t. ‘Clarify yes-no’. The worst performance in terms of KE and QS is achieved w.r.t. ‘Clarify yes-no’ and in terms of AP and RG w.r.t. ‘Clarify open’. For response generation w.r.t. ‘Clarify yes-no’, WISE can almost copy the answer from a single query, which is easier than response generation w.r.t. ‘Clarify choice’ and ‘Clarify open’, which different queries, like ‘city’ from ‘Beijing’ and ‘Jinan’.

Second, for the action ‘Answer’, the best performance in terms of KE and AP is achieved w.r.t. ‘Answer-open’ and in terms of PS and RG w.r.t. ‘Answer-fact’. The worst performance in terms of all tasks is achieved w.r.t. ‘Answer-opinion’. For response generation w.r.t. ‘Answer fact’, WISE can copy rather than summarize from multiple passages w.r.t. ‘Answer open’ and ‘Answer opinion’.

Third, the performance of PS has a high correlation with the final RG task w.r.t. ‘Answer’, while the performance of QS does not seem to be correlated w.r.t. ‘Clarify’. As we can see in Table 8, an F1 score of 23.3% in PS and BLEU score of 12.6% in terms of RG is achieved w.r.t. ‘Answer opinion’, and a higher F1 score (30.0%) in terms of PS and a higher BLEU score (16.4%) in terms of RG w.r.t. ‘Answer fact’. For response generation w.r.t. ‘Answer’, large portions of the output can be copied from passages, and thus the performance PS is crucial to the final RG performance. The F1 score of WISE is 10.9% in QS w.r.t. ‘Clarify yes-no’, which is lower than 29.6% w.r.t. ‘Clarify choice’; the performance in RG is better, e.g., 28% vs. 20%. It is more likely that WISE has to copy from multiple queries for ‘Clarify choice’ and from a single query for ‘Clarify yes-no’, leading to performance differences for QS and RG, w.r.t. ‘Clarify’.

We split the test dataset based on the occurrence of conversational background in the training dataset: (test (unseen) means no occurrence, and test (seen) means occurrence); the results are shown in Table 7. Generally, the performances of WISE on all tasks is best on the test (seen) dataset, and worst on the test (unseen) dataset. For example, it achieves 29.6% in terms of F1 for PS in test (seen) dataset, which is higher than 24.9% in test dataset and 20.5% in test (unseen) dataset. WISE has been trained on the same conversational background in the test (seen) condition, and has never seen the conversational background in the test (unseen) condition.

### 6.2 Effects of subtasks

We analyze the effects of different tasks in WISE, w.r.t. ID, KE, AP, QS and PS, and the results are shown in Table 9. We remove each task, resulting in five settings; ‘-xx’ means WISE without module ‘xx’. Generally, removing any module will result in a drop in performance for RG on all metrics, indicating that all parts are helpful to WISE. Without QS, RG drops sharply on all metrics. Specially, it

**Table 6: Overall performance. Bold face indicates the best result in terms of corresponding metric. Significant improvements over the strongest WISE results are marked with \* (p-value < 0.01 with t-test).**

	ID (%)			KE (%)		AP (%)			QS (%)			PS (%)			RG (%)	
	P	R	F1	BLEU	ROUGE	P	R	F1	P	R	F1	P	R	F1	BLEU	ROUGE
WISE-pretrain	<b>71.3</b>	12.5	11.7	0.0	0.0	12.7	4.5	1.8	0.0	0.0	0.0	0.0	0.0	0.0	1.0	1.0
WISE	45.2	<b>32.5</b>	<b>34.1</b>	<b>58.3</b>	<b>59.0</b>	<b>18.8</b>	<b>20.6</b>	<b>17.8</b>	<b>26.9</b>	<b>16.1</b>	<b>11.5</b>	<b>35.8</b>	<b>31.5</b>	<b>24.9</b>	14.6	16.9
WISE+GT	45.2	<b>32.5</b>	<b>34.1</b>	<b>58.3</b>	<b>59.0</b>	<b>18.8</b>	<b>20.6</b>	<b>17.8</b>	<b>26.9</b>	<b>16.1</b>	<b>11.5</b>	<b>35.8</b>	<b>31.5</b>	<b>24.9</b>	<b>20.4*</b>	<b>24.2*</b>

**Table 7: Overall performance of WISE on the test(seen) and test(unseen) dataset.**

	ID (%)			KE (%)		AP (%)			QS (%)			PS (%)			RG (%)	
	P	R	F1	BLEU	ROUGE	P	R	F1	P	R	F1	P	R	F1	BLEU	ROUGE
test (unseen)	38.4	28.5	29.3	<b>59.2</b>	<b>59.0</b>	17.6	18.1	16.5	<b>28.8</b>	15.1	10.8	34.2	26.7	20.5	14.1	16.0
test (seen)	<b>48.3</b>	<b>36.1</b>	<b>37.4</b>	57.3	58.4	<b>19.9</b>	<b>24.2</b>	<b>19.0</b>	25.4	<b>21.5</b>	<b>15.4</b>	<b>36.8</b>	<b>37.6</b>	<b>29.6</b>	<b>14.6</b>	<b>17.4</b>
test	45.2	32.5	34.1	58.3	<b>59.0</b>	18.8	20.6	17.8	26.9	16.1	11.5	35.8	31.5	24.9	<b>14.6</b>	16.9

**Table 8: Overall performance of WISE for different actions.**

		KE (%)		AP (%)			QS (%)			PS (%)			RG (%)	
		BLEU	ROUGE	P	R	F1	P	R	F1	P	R	F1	BLEU	ROUGE
Clarify	yes-no	65.3	67.9	<b>46.9</b>	15.0	22.7	22.4	20.0	10.9	—	—	—	<b>28.0</b>	<b>30.7</b>
	choice	<b>74.3</b>	<b>75.9</b>	22.4	<b>37.4</b>	28.0	<b>46.4</b>	<b>38.0</b>	<b>29.6</b>	—	—	—	20.0	21.7
	open	67.2	68.9	8.4	17.6	11.4	38.9	22.1	16.7	—	—	—	15.0	16.3
Answer	opinion	57.1	57.2	2.3	17.3	3.2	—	—	—	32.3	<b>37.0</b>	23.3	12.6	13.6
	fact	60.1	61.2	13.5	21.2	15.7	—	—	—	<b>44.1</b>	34.6	<b>30.0</b>	16.4	19.6
	open	62.7	62.8	39.1	35.6	<b>36.7</b>	—	—	—	41.6	35.1	29.1	13.4	17.2

drops 4.2% in terms of BLEU for RG, which means QS is essential to RG. While without PS, RG drops a little compared to QS. Specially, it drops only 0.9% in terms of BLEU for RG. The performance of PS is far from perfect, only 24.9% in F1 score; omitting it has limited impact. In addition, removing ID, KE or AP harms the performance of QS and PS, indicating that all previous modules for QS and PS help for the performance of QS and PS. Specifically, performance drops by 10% and 5.7% in terms of F1 score for QS and PS when removing KE: if the model does not extract the accurate keyphrase, it can not select the most relevant queries and passages.

### 6.3 Effects of pretraining

To analyze the effects of different pretrain datasets, we conduct an ablation study (Table 10). WISE uses four external datasets for pretraining; each time we remove one, we use the others for pretraining. Hence, we have four conditions. In Table 10, ‘-xx’ means pretraining WISE using datasets other than ‘xx’.

The results show that all datasets are helpful for pretraining; removing any dataset results in a performance drop for RG. The results drop in terms of all metrics when pretraining without the DuReader dataset. The results drop by 5.0% and 5.9% in terms of BLEU and ROUGE for RG. The drop after removing pretraining with DuReader is far larger than after removing any of the other individual pretraining dataset, indicating the importance of DuReader for pretraining, as shown in Table 5. Interestingly, sometimes removing a dataset will improve the performance of ID and AP. For example, -duconv achieves 35.8% in terms of F1 in ID, higher than

the 34.1% for WISE. ID and AP predict labels among small sets, which are relatively easy compared to the other tasks.

## 7 RELATED WORK

*User studies.* To determine whether CIS is needed and what it should look like, researchers have conducted different user studies. Vtyurina et al. [42] 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 conclude that people do not have biases against CIS systems, as long as their performance is acceptable. Trippas et al. [37] conduct a laboratory-based observational study, where pairs of people perform search tasks communicating verbally. They conclude that CIS is more complex and interactive than traditional search. Based on their conclusions, we can confirm that CIS is the right direction but there is still a long way to go.

*Datasets.* Several conversation related datasets are available, for various tasks, e.g., knowledge grounded conversations [14, 22], conversational question answering [8, 10, 29, 32]. However, their targets are different from information seeking in the search scenarios.

To address this gap, there are various efforts to build datasets for CIS. Dalton et al. [12] release the CAsT dataset, which aims to establish a concrete and standard collection of data to make CIS systems directly comparable. Ren et al. [31] provide natural



**Table 9: Overall performance of models removing different tasks.**

	ID (%)			KE (%)		AP (%)			QS (%)			PS (%)			RG (%)	
	P	R	F1	BLEU	ROUGE	P	R	F1	P	R	F1	P	R	F1	BLEU	ROUGE
-ID	—	—	—	54.6	54.8	18.7	22.6	18.3	25.5	8.3	6.1	<b>43.5</b>	25.1	22.7	13.0	15.2
-KE	<b>66.2</b>	28.2	30.6	—	—	22.0	22.7	<b>19.1</b>	<b>28.0</b>	1.6	1.5	39.8	20.8	19.2	12.6	14.7
-AP	52.5	32.2	<b>35.3</b>	<b>59.5</b>	<b>59.4</b>	—	—	—	21.6	11.0	8.5	37.8	26.4	21.9	13.2	15.1
-QS	51.9	<b>32.6</b>	32.6	54.2	54.2	<b>22.7</b>	<b>23.2</b>	18.9	—	—	—	39.5	25.9	22.1	10.4	12.1
-PS	51.3	30.8	32.8	56.8	57.0	20.1	22.3	18.1	24.7	8.1	6.3	—	—	—	13.7	15.7
WISE	45.2	32.5	34.1	58.3	59.0	18.8	20.6	17.8	26.9	<b>16.1</b>	<b>11.5</b>	35.8	<b>31.5</b>	<b>24.9</b>	<b>14.6</b>	<b>16.9</b>

**Table 10: Overall performance of WISE removing different datasets for pretraining.**

	ID (%)			KE (%)		AP (%)			QS (%)			PS (%)			RG (%)	
	P	R	F1	BLEU	ROUGE	P	R	F1	P	R	F1	P	R	F1	BLEU	ROUGE
-DuReader	<b>47.5</b>	25.7	27.7	50.7	51.8	18.0	19.5	17.2	<b>33.3</b>	2.2	2.4	33.7	17.5	15.1	9.6	11.0
-KdConv	41.1	27.7	28.1	49.7	51.3	16.3	17.7	15.3	22.2	8.4	6.3	<b>37.2</b>	26.0	22.8	12.5	14.7
-DuConv	43.9	<b>35.5</b>	<b>35.8</b>	55.8	57.0	20.3	20.2	17.9	30.8	7.7	6.6	34.8	28.0	22.6	13.4	15.4
-WebQA	39.0	30.6	32.0	57.1	57.4	<b>20.9</b>	<b>20.9</b>	<b>18.8</b>	27.7	10.6	8.4	34.9	23.6	19.6	12.3	14.3
WISE	45.2	32.5	34.1	<b>58.3</b>	<b>59.0</b>	18.8	20.6	17.8	26.9	<b>16.1</b>	<b>11.5</b>	35.8	<b>31.5</b>	<b>24.9</b>	<b>14.6</b>	<b>16.9</b>

language responses by summarizing relevant information in paragraphs based on CAsT. Aliannejadi et al. [2] explore how to ask clarifying questions, which is important in order to achieve mixed-initiative CIS. They propose an offline evaluation methodology for the task and collect the dataset ClariQ through crowdsourcing. The above datasets only cover some aspects of CIS.

To better understand CIS conversations, Thomas et al. [35] release the MISC dataset, which is built by recording pairs of a seeker and intermediary collaborating on CIS. The work most closely related to ours is by Trippas et al. [38]. They create the CSCData dataset for spoken conversational search (very close to the CIS setting), define a labelling set identifying the intents and actions of seekers and intermediaries, and conduct a deep analysis. The MISC and CSCData datasets only have 666 and 1,044 utterances, which results in a low coverage of search intents and answer types. No supporting documents are available and the responses are in very verbose spoken language. Importantly, it is hard for the defined labels to make a direct connection to TSEs.

*Frameworks.* Several recent frameworks address a specific aspect of CIS, e.g., asking clarifying questions [48], generating responses for natural language questions [23]. There are also attempts to design a complete CIS that could work in practice. Radlinski and Craswell [28] consider the question of what properties would be desirable for a CIS system so that the system enables users to answer a variety of information need in a natural and efficient manner. They present a theoretical framework of information interaction in a chat setting for CIS, which gives guidelines for designing a practical CIS system. Azzopardi et al. [4] outline the actions and intents of users and systems explaining how these actions enable users to explore the search space and resolve their information need. Their work provides a conceptualization of the CIS process and a framework for the development of practical CIS systems.

To sum up, the work listed above either focuses on a particular aspect or studies the modeling of CIS theoretically.

## 8 CONCLUSION AND FUTURE WORK

We have proposed a pipeline and a modular end-to-end neural architecture for conversational information seeking (CIS), which models CIS as six sub-tasks to improve the performance on the response generation (RG) task. We also collected a WISE dataset in a wizard-of-oz fashion, which is more suitable and challenging for comprehensive and in-depth research on all aspects of CIS. Extensive experiments on the WISE dataset show that the proposed WISE model can achieve state-of-the-art performance and each proposed module of WISE contributes to the final response generation.

Although we have proposed a standard dataset for CIS, we have a long way to go to improve the performance on each subtask, and to consider more factors, e.g., the information timeliness [15]. A potential direction is to leverage large-scale unlabeled datasets on related tasks like dialogue and information retrieval.

## REPRODUCIBILITY

The code and data used to produce the results in this paper are available at [https://github.com/PengjieRen/Case\\_WISE](https://github.com/PengjieRen/Case_WISE).

## ACKNOWLEDGMENTS

We thank the reviewers for their valuable feedback. This research was partially supported by the National Key R&D Program of China with grant No. 2020YFB1406704, the Natural Science Foundation of China (61972234, 61902219, 62072279), the Key Scientific and Technological Innovation Program of Shandong Province (2019JZZY010129), the Tencent WeChat Rhino-Bird Focused Research Program (JR-WXG-2021411), the Fundamental Research Funds of Shandong University, and the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research, <https://hybrid-intelligence-centre.nl>. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

## REFERENCES

- [1] Abien Fred Agarap. 2018. Deep Learning using Rectified Linear Units (ReLU). *arXiv preprint arXiv:1803.08375* (2018).
- [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 and Development in Information Retrieval, SIGIR 2019*. 475–484.
- [3] Avishek Anand, Lawrence Cavedon, Matthias Hagen, Hideo Joho, Mark Sanderson, and Benno Stein. 2020. Conversational Search - A Report from Dagstuhl Seminar 19461. *arXiv preprint arXiv:2005.08658* (2020).
- [4] 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*.
- [5] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [6] Ashutosh Baheti, Alan Ritter, and Kevin Small. 2020. Fluent Response Generation for Conversational Question Answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*. 191–207.
- [7] Andrei Z. Broder. 2002. A taxonomy of web search. *SIGIR Forum* 36 (2002), 3–10.
- [8] 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 2018*. 2174–2184.
- [9] 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, SIGKDD 2016*. 815–824.
- [10] 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 2019*. 729–738.
- [11] Adam Czyzewski, Jeffrey Dalton, and Anton Leuski. 2020. Agent Dialogue: A Platform for Conversational Information Seeking Experimentation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020*. 2121–2124.
- [12] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. TREC CAsT 2019: The Conversational Assistance Track Overview. *arXiv preprint arXiv:2003.13624* (2020).
- [13] 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 2020*. 1985–1988.
- [14] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-Powered Conversational Agents. In *Proceedings of the 7th International Conference on Learning Representations, ICLR 2019*.
- [15] Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. Temporal Relational Ranking for Stock Prediction. *ACM Trans. Inf. Syst.* 37, 2 (2019), 30 pages.
- [16] Chongming Gao, Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2021. Advances and Challenges in Conversational Recommender Systems: A Survey. *arXiv preprint arXiv:2101.09459* (2021).
- [17] Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu, Qiaoqiao She, Xuan Liu, Tian Wu, and Haifeng Wang. 2018. DuReader: a Chinese Machine Reading Comprehension Dataset from Real-world Applications. In *Proceedings of the Workshop on Machine Reading for Question Answering, ACL 2018*. 37–46.
- [18] Tomas Kocisky, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gabor Melis, and Edward Grefenstette. 2018. The NarrativeQA Reading Comprehension Challenge. *Trans. Assoc. Comput. Linguistics* 6 (2018), 317–328.
- [19] Bernhard Kratzwald and Stefan Feuerriegel. 2019. Learning from On-Line User Feedback in Neural Question Answering on the Web. In *Proceedings of the World Wide Web Conference, WWW 2019*. 906–916.
- [20] Peng Li, W. Li, Z. He, Xuguang Wang, Y. Cao, J. Zhou, and W. Xu. 2016. Dataset and Neural Recurrent Sequence Labeling Model for Open-Domain Factoid Question Answering. *arXiv preprint arXiv:1607.06275* (2016).
- [21] Chin-Yew Lin. 2004. Rouge: A Package for Automatic Evaluation of Summaries. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics, ACL 2002*. 74–81.
- [22] Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. Towards Exploiting Background Knowledge for Building Conversation Systems. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, EMNLP 2018*. 2322–2332.
- [23] 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 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems NIPS 2016*, Vol. 1773.
- [24] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, ACL 2002*. 311–318.
- [25] Jiahuan Pei, Pengjie Ren, Christof Monz, and Maarten de Rijke. 2020. Retrospective and Prospective Mixture-of-Generators for Task-Oriented Dialogue Response Generation. In *Proceedings of the 24th European Conference on Artificial Intelligence, ECAI 2020*, Vol. 325. 2148–2155.
- [26] Baolin Peng, Xiujun Li, Lihong Li, Jianfeng Gao, Asli Celikyilmaz, Sungjin Lee, and Kam-Fai Wong. 2017. Composite Task-Completion Dialogue Policy Learning via Hierarchical Deep Reinforcement Learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017*. 2231–2240.
- [27] Chen Qu, Liu Yang, W. Bruce Croft, Yongfeng Zhang, Johanne R. Trippas, and Minghui Qiu. 2019. User Intent Prediction in Information-seeking Conversations. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval, CHIIR 2019*. 25–33.
- [28] 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 2017*. 117–126.
- [29] Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Trans. Assoc. Comput. Linguistics* 7 (2019), 249–266.
- [30] 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 Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*. 8697–8704.
- [31] Pengjie Ren, Zhumin Chen, Zhaochun Ren, Evangelos Kanoulas, Christof Monz, and Maarten de Rijke. 2020. Conversations with Search Engines: SERP-based Conversational Response Generation. *arXiv preprint arXiv:2004.14162* (2020).
- [32] 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, AAAI 2018*. 705–713.
- [33] 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 2017*. 1073–1083.
- [34] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2018*. 235–244.
- [35] Paul Thomas, Mary Czerwinski, Daniel J. 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 2018*. 42–51.
- [36] 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 and Retrieval, CHIIR 2018*. 32–41.
- [37] 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 and Retrieval, CHIIR 2018*. 32–41.
- [38] Johanne R. Trippas, Damiano Spina, Paul Thomas, Mark Sanderson, Hideo Joho, and Lawrence Cavedon. 2020. Towards a Model for Spoken Conversational Search. *Inf. Process. Manag.* 57 (2020), 102162.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS 2017*. 5998–6008.
- [40] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer Networks. In *Proceedings of 28th Annual Conference on Neural Information Processing Systems, NIPS, 2015*. 2692–2700.
- [41] 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 2020*. 921–930.
- [42] Alexandra Vtuyrina, 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 on Human Factors in Computing Systems, CHI, 2017*. 2187–2193.
- [43] Ryan W. White and Resa A. Roth. 2009. *Exploratory Search: Beyond the Query-Response Paradigm*.
- [44] Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive Human-Machine Conversation with

- Explicit Conversation Goal. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*. 3794–3804.
- [45] Wei Wu and Rui Yan. 2019. Deep Chit-Chat: Deep Learning for Chatbots. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019*. 1413–1414.
- [46] Zhijing Wu, Jiaxin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2019. Investigating Passage-level Relevance and Its Role in Document-level Relevance Judgment. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019*. 605–614.
- [47] Hamed Zamani and Nick Craswell. 2020. Macaw: An Extensible Conversational Information Seeking Platform. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020*. 2193–2196.
- [48] Hamed Zamani, Susan T. Dumais, Nick Craswell, Paul N. Bennett, and Gord Lueck. 2020. Generating Clarifying Questions for Information Retrieval. In *Proceedings of the Web Conference, WWW 2020*. 418–428.
- [49] Hao Zhou, Chujie Zheng, Kaili Huang, Minlie Huang, and Xiaoyan Zhu. 2020. Kd-Conv: A Chinese Multi-domain Dialogue Dataset Towards Multi-turn Knowledge-driven Conversation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*. 7098–7108.