

Conversational Interfaces for Information Search

Q. Vera Liao, Werner Geyer, Michael Muller, Yasaman Khazaen

Abstract Recent progress in Machine Learning has given rise to a plethora of tools and applications that rely on conversational interactions, from chatbots, speech-controlled devices to robots and virtual agents. Conversational interfaces are becoming widely accepted for utility tools, where a common function is to serve users' information needs. Albeit with much excitement, we are only starting to understand how users' information seeking behaviors and design opportunities may transform moving from traditional graphical user interfaces to conversational user interfaces. In this chapter, we start by reviewing recent work in the emerging area of conversational interfaces, and lay out their opportunities for supporting information search tasks. We then present insights from our experience deploying a chatbot supporting information search in a large enterprise, demonstrating how a conversational interface impacts user behaviors and offers new opportunities for improving search experience, in particular for user modeling.

1 Introduction

Conversational systems, with which users interact through spoken or typed natural language, are becoming ubiquitous thanks to the popularity of many commercial products such as Apple Siri, Google Assistant, and Amazon Alexa. We *talk* to these systems to find directions, choose products, troubleshoot technical problems, learn about the latest news and seek information for various decisions. This paradigm shift to conversational interactions on the one hand offers new opportunities to improve

Q. Vera Liao
IBM Research AI, Yorktown Heights, NY, USA e-mail: vera.liao@email.address

Werner Geyer, Michael Muller, Yasaman Khazaen
BM Research AI, Cambridge, MA, USA e-mail: {werner.geyer, michael_endmathmuller, yasaman.khazaeni}@us.ibm.com

information-seeking tasks, on the other hand, requires developing new types of computational models and interaction techniques.

If a conversational system is seen as merely a text-based or voice-based search system, it makes sense to use search results to respond to user questions. For example, the following exchange is a common scenario in talking to a FAQ chatbot on a web page of an insurance company:

User: Is it mandatory to purchase collision coverage?

Agent: Collision coverage pays for repair or replacement costs if your car crashes into an object or another car. This coverage is designed to fix up or replace your own car after a collision. Our customers will need to have comprehensive coverage before adding collision coverage. You can have comprehensive and not collision on your policy, just not the other way around. If you lease or finance your car, collision may be required by your leasing or financing company. Like comprehensive coverage, collision coverage will pay for damages to your car as long as they don't exceed the car's actual cash value...

However, such an interaction is not yet compatible with how humans engage in conversations. Much burden is still on the user to locate the answer from the lengthy texts of the search result. More problematically, the customer has to accurately formulate the question or the system would retrieve an irrelevant set of results. In that case, the customer would have to carefully read the large chunk of text, decide there is no useful information, and start over again by reformulating the question, leading to frustration or boredom. We envision that one day we could actually talk to intelligent systems as follows:

User: Hi agent, I need some help

Agent: Sure! How can I help you?

User: Do I need to buy coverage for my own car?

Agent: Are you considering collision coverage or comprehensive coverage?

User: What is collision coverage?

Agent: If your car crashes into an object or another car, collision coverage is designed to fix up or replace your own car.

User: How is it different from comprehensive coverage?

Agent: Well, collision coverage is to cover your own car after a collision, while comprehensive coverage pays for damage in events that are not collision-related, such as theft and vandalism.

User: I see. Am I required to purchase collision coverage then?

Agent: Collision coverage is optional. But if you lease or finance your car, it may be required by your leasing or financing company

User: Okay thanks.

Agent: You are welcome!

In such conversational interactions, information is given by the system in a succinct and precise manner. More importantly, an accurate understanding of the customer's problem is incrementally formed through various actions from the system such as clarification and explanation, which are made possible by having the customer well engaged in a natural and personable conversation. Conversational interface like

this can reduce customers' cognitive workload to read, requires little learning to use, and might be the only viable solution in some contexts such as voice-based systems.

The agents in both examples have access to the same set of information sources. So how do we move from single-turn Q&A in Example 1 to multi-turn interactions in Example 2? By talking naturally, users should be able to express their information needs in a familiar way as in their everyday conversations, no longer required to learn the operation vocabulary of a search system or the output schema. This means that users would be offloaded from "direct manipulation" of computational models [52], instead relying on conversational interfaces to perform an *intermediary* role between the information seekers and information sources. Natural conversation offers the necessary devices to perform such an intermediary role, as we communicate through language to exchange information between our different minds, not through queries-and-answers, but non-linear combinations of disclosure, elicitation, refinement, clarification, explanation, narratives and so on. Conversational interfaces should take advantage of these devices

This is not to say that conversational systems should mimic entire human conversations, nor could we expect them to achieve human-level intelligence any time soon. Rather, our view is that we should start with the fundamental question—what aspects of conversations are pertinent and beneficial for performing information search tasks, and what are the necessary system capabilities to enable them?

We believe these questions should be addressed from two ends: From a systems point of view, i.e. how can conversational interactions extend current computational models for information search; From a cognitive point of view, i.e. what are the desirable and necessary properties of conversations to support information seeking. In the next section, we offer some perspectives by drawing on related work that answers the following questions: 1) How do conversational interactions fit and extend information search models; 2) Empirically, what have we learned about designing functionalities of conversational search systems, which are still very much in an exploratory stage; 3) Theoretically, what are the fundamental properties of human conversations, and how can they inform the design of conversational search to make it more cognitively compatible with how people converse.

With that, we lay out a space of functional goals for conversational search systems along two axes: stages of information search behaviors—query formulation, search result exploration, and repair [6, 42, 63]; and fundamental principles of human conversations—efficiency, common ground and recipient design [9, 21, 35, 50]. We also discuss empirical insights on how users converse with a conversational agent performing information search tasks from our own work deploying a chatbot that answers questions from employees of a large enterprise. The results demonstrate the rich conversational behaviors users engage with a conversational interface, and the opportunities they offer for improving search performance and user experience.

2 Conversational Search

Conversational systems that support information-seeking tasks encompass systems with a one-shot question-and-answer (QA) model (as in most current commercial products), systems that can engage in free conversations to resolve a user’s problem, and anything in between. To be excluded are systems that perform primarily chit-chat, and systems that receive or give commands through voice or text. This chapter will not discuss 60-years history of dialogue systems and conversational agents. For interested readers we recommend several books [11, 32, 33]. We do not address with any particular computational solutions, whether rule-based, statistical, or more recent neural network models [7, 38, 55]. Our focus is on designing *interactions* of conversational search.

While people have long been fascinated by “machines to talk to”, migrating to a new interaction modality requires work to define new system functionalities and design guidelines. The information retrieval (IR) community responded with much enthusiasm—“*such a growth in natural language dialogue between users and search systems may even lead to the dominant interaction model of one-shot keyword queries being displaced with conversational systems*” [49]. The IR community takes a system-driven approach and considers new opportunities offered by conversational interactions to extend existing IR frameworks [3, 14, 38, 49, 67]. The human-computer interaction (HCI) community, expressing equal enthusiasm, follows its user-centered design tradition by studying users’ needs and behaviors to inform the design of this new type of interface [40, 41, 46, 48]. We draw on work from both communities as well as related social and cognitive science to reflect on how to best deliver the benefits of conversational interactions for information search tasks.

2.1 From Search to Conversational Search

Across academic communities, there is a long history of creating formal frameworks of users’ information-seeking behaviors to guide IR system design [6, 42, 63]. While details vary, these frameworks converge on three core stages of information-seeking: query formulation from information needs, examination of search results, and query refinements or reformulation if necessary. Chapter 13 of this book provides a more detailed overview of these multi-stage search behavioral models. Earlier work sought inspiration from structures of information-seeking dialogues (e.g., with a librarian) to inform *interactive search*. Influential work includes Conversation Role Model by Sitter and Stein [54] and conversational scripts by Belkin et al. [8]. Both aimed to provide idealized flows for an interactive search system to select next steps, but they were not considered specifically for conversational systems.

Recently, researchers started conceptualizing what actions can be made by users and systems if search becomes “conversational”. In a perspective paper by Radlinski and Craswell [49], they identify properties of conversations pertinent to search settings to be: mixed-initiative, goal oriented, maintaining memory, and adaptive.

USER			AGENT		
	Query Formulation	Reveal Disclose Revise Refine Expand	Inquire Extract Elicit Clarify	<i>User Disclosure</i>	
<i>Set Retrieval</i>	Result Exploration	Inquire List Summarize Compare Subset Similar	Reveal List Summarize Compare Subset Similar	<i>System disclosure</i>	<i>Memory</i>
		Navigate Repeat Back More Note	Traverse Repeat Back More Record		
<i>Mixed Initiative</i>		Interrupt Interrupt	Suggest Recommend Hypothesize		
		Interrogate Understand Explain	Explain Report Reason		

Fig. 1 An action space of conversational search systems proposed by [3] based on properties of conversational search proposed by [49]

Accordingly, they offer a formal definition of conversational search system to be “a system for retrieving information that permits a mixed-initiative back and forth between a user and agent, where the agent’s actions are chosen in response to a model of current user needs within the current conversation, using both short- and long-term knowledge of the user.

Radlinkski and Craswell argue that a conversational search system should therefore have the following properties as advantages over traditional search systems : 1) User disclosure: through conversations, it could help the user better express information needs; 2) System disclosure: it is convenient for the system to reveal its capabilities, building the user’s mental model; 3) Mixed-initiative: the system and user can both naturally take initiative as appropriate; 4) Memory: the user can naturally reference past statements; 5) Set retrieval: the system can reason about the utility of a set of items in a conversation rather than having multiple search sessions as in traditional search system. Based on these requirements, Azzopardi et al. [3] proposed possible actions that a user and a conversational agent could perform during query formulation and search results exploration stages, as well as mixed-initiative actions that can happen in the conversation (Figure 1). This action spaces could be used to generate dialogue policies for performing conversational search.

To complement these theoretical perspectives, in the following, we discuss empirical work on developing functionalities of conversational search. We will organize the discussions by the three core stages of information-seeking behaviors. Currently, to handle free-form conversations remains an open challenge. Most work addresses

only some aspects of conversational search. Another approach taken is to study how people naturally engage in information-seeking conversations [43, 60], either with another human, or a wizard-of-oz agent (an unseen human simulating an agent). The premise is that these scenarios represent the ideal level of intelligence and that users prefer interactions consistent with their natural conversations. While such a view is debatable, it is a valuable approach to seek inspirations to define capabilities and actions of conversational search.

2.1.1 Query formulation

Whether through a speech-controlled device or chatting with a chatbot, users' querying behaviors may naturally change when the interface becomes conversational. A number of studies examined user queries with spoken search systems, and compared them to typed queries in a search box, showing that spoken queries are longer, more verbose, and have more varied language [16, 23]. For example, Guy conducted an analysis on half a million search logs and found that spoken queries have more formal grammatical structures (e.g., wh-words—what/why/who/where), types of part-of-speech (while typed queries are mostly nouns), and tend to use full-sentence inquiries (e.g., "I am looking for", "take me to") [23]. Trippas et al. conducted a study observing conversations between an information seeker and a human intermediary who had access to a search system [60]. They observed much variance—while some used query-like expressions, others used lengthy and complex sentences to describe their needs. These natural-language expressions could include multiple actions (e.g., querying, navigation) in one turn, or complete one action with multiple turns, in contrast to a linear process with a traditional search interface.

Perhaps the most critical difference that a conversational system makes is that query formulation can be achieved through multi-turn interactions [15, 30, 59, 67]. It means that, for one, it is possible for the user to express complex information needs sequentially, yielding a more accurate representation of their real knowledge gap, which may be multi-faceted, multi-item, or inter-dependent [49]. Second, the system could take an active role that resembles an intermediary, by asking questions back to the user depending on the previous information provided, thus offloading the user effort to accurately formulate formal system input.

According to Azzopardi's model [3], during the query formulation stage, a user may follow up by *revising*, *refining* and *expanding* the original query, and a system can "*extract, elicit and clarify*." Extracting key information and clarification are necessary actions to deal with the verbose, sometimes ambiguous queries in conversational forms. Eliciting user criteria, preferences, or constraints to refine search results has long been of interest to IR systems [5, 13, 31]. Earlier work on goal-oriented dialogue systems also adopted a "slot-filling" approach that requests one criterion from the user at a time [10, 61]. The outcome, however, might be a long, tedious dialogue that does not feel natural [68]. This approach is also questionable outside narrow domains if the criteria have many candidates or cannot be pre-defined. An alternative approach is to acquire user criteria by eliciting feed-

back for sample items through critiquing [31], comparing [14], or grouping similar items [13]. By suggesting items and eliciting feedback, it could also help build users' mental models of the search space. These approaches could face challenges with traditional search system as users may be unwilling to provide feedback by repeatedly filling out forms. Conversational interfaces could be a natural fit, as example-based discussions are common in conversations, and with natural language the user input can be incremental and flexible.

There are a number of computational challenges to enable query formulation through conversations. First, extraction and pre-processing techniques are necessary to bridge natural language queries and input for the underlying computational models. Currently, to handle variations in conversational input, rule-based or machine learning based intent models are often used to first map a user utterance to a query category for system input. Argullo et al. showed that conversational queries yield worse retrieval performance if issued unmodified to search APIs [1]. While Crestani and Du suggested that simple processing by extracting nouns, adjectives and verbs can improve the retrieval performance [16], the challenges are likely beyond keywords extractions. For example, one may need to breakdown multiple moves in one utterance, understand user needs from multiple turns, and accommodate greater individual variances in querying behaviors. Second, optimizing the dialogue flow for elicitation questions remains a computational challenge. To efficiently reduce the search space (sometimes relaxing if over-specified) involves multiple complex optimization problems, such as selecting elicitation methods to use, criteria to elicit for, and considering the trade-off between further elicitation and revealing results.

2.1.2 Search results presentation

How to present search results in conversational forms is an under-studied area. Most existing systems simply either read out the top result or a condensed version of a ranked list. Such designs are not conversational, and can be problematic for speech generation and users' attention span (for listening and reading). Also many key elements of traditional search systems, such as search snippets and graphic information, cannot be easily presented in conversational forms. It is important to recognize that, compared to graphical user interfaces, conversations represent a "narrower" and thus ideally more precise information channel—a turn of a dialogue is expected to be succinct, informative and relevant [21]. Therefore additional processing mechanisms are required on top of document retrieval.

One such mechanism is summarization. Trippas et al.'s study shows that the human intermediary naturally chose to provide document summarization to the information seekers, sometimes summarizing across multiple documents [60]. The counterparts of summarization—drilling-down for details—would also be necessary, which are embodied in various actions in natural conversations, mainly paraphrasing, defining, explanation, and elaboration [51]. Azzopardi et al.'s model emphasizes the presentation of multiple documents, suggesting that users should be able to request presentation of multi-documents in overview—*summarization*, *listing*, etc.,

or in details—*subsets, comparison between documents, etc.*, and to navigate between documents—*forward, back-ward and repeat* [3].

The technical challenge here is to create representations of varied granularity for a document or a set of documents, from high-level summarization to different types of details. In their book on design patterns for conversational systems, Moore and Arar suggested patterns to decompose document content, such as a FAQ page, for conversational interactions [45]. While a search engine may simply retrieve a document (e.g. health insurance coverage), a conversational system should not only recognize sub-topics within the document (e.g., copay, deductible), but also respond to different types of follow-up requests—paraphrasing, examples and definitions, by either generating or extracting such contents from the document.

Conversational interfaces are especially suitable for driving the user down a focused navigational path. The drawback is that the user may lose awareness of alternate information, and risk forming a narrow understanding or inaccurate mental model of the search space. This is another place where the system should take initiative. Azzopardi et al. suggest that the system should always keep a representation of users' current information needs, past information needs, as well as alternate information needs, and provide recommendations that the user may not have explicitly requested [3]. However, technical challenges remain to be solved, not only on how to infer alternate information needs, but also when to suggest them. For example, one needs to consider the trade-off between access to more information and the increasing complexity of the conversation.

2.1.3 Query reformulation and repair

In search behavior models, query reformulation is the step following dissatisfying search results exploration. This definition is worth revisiting in a conversational search setting. On the one hand, in conversational search the querying process can be incremental, i.e. query formulation and search result presentation can happen multiple times in series (e.g., to elicit further feedback), thus blurring the boundaries between query formulation and reformulation. On the other hand, conversational systems introduce new types of errors in addition to sub-optimal retrieval performance, such as errors in the steps of speech recognition, speech to text, or language understanding. Hence, it may be more appropriate to use the term “repair” to consider user and system actions when there is breakdown in conversational search.

Recent work investigated how users naturally reformulate queries with conversational search systems [24, 26, 46, 53]. For example, with a spoken search system, Jiang et al. showed that users engaged in various ways of lexical reformulation (e.g., addition, substitution, removing, reordering) and phonetic re-formulation (e.g. emphasizing parts of the query) [26]. However, it creates problem when users make blind attempts of reformulation, which sometimes contradicts with practices in natural conversations. For instance, while it is natural for people to elaborate with more information when talking to another person, reformulating queries with more details may adversely harm the retrieval performance [46]. This is due to a mismatch

in the natural language understanding capabilities of a person versus a computational system. Users often do not understand this mismatch because conversational interfaces are misleadingly “natural” and thus opaque in disclosing the underlying computational models.

A system may fail both the search task and creating a truly conversational experience if solely relying on users to repair, as conversation should be a two-way process to resolve uncertainty [35]. Recent work started addressing system-initiated repairing processes [4, 47]. In [2], the authors propose a framework for repair design of conversational QA with three levels of increasing contributions from the system: explicitly acknowledging the breakdown (e.g., asking for confirmation or rephrasing), making the system model transparent to assist user repair (e.g., explaining current understanding), and proactively suggesting query reformulation. The research shows that user satisfaction increases with a higher level of contribution from the system.

The technical challenge for the three levels of system-initiated contributions, however, remains largely unsolved. First, there is often no precise way for a system to recognize a misunderstanding or retrieval error. Currently, machine-learning based systems rely on a confidence level to infer a potential breakdown. Such methods cannot identify “unknown unknown” errors where the training data have blind spots. Paek and Horvitz explored using Bayesian networks to infer uncertainty [47], but it has not been adopted as a scalable solution. Second, explaining language understanding or search results is a pressing problem that has raised much interest lately [22, 44]. The challenge is on how to make the explanation actionable to support the end goal of successful query reformulation. For example, a system can explain its current understanding by keywords it identified for search results, and the user should be able to respond by incrementally correcting misunderstanding without starting all over. Lastly, as with traditional search systems, query suggestions may fail, and such failures are likely more detrimental if presented as a single turn in a conversation.

2.2 Towards Conversational User Experience

Now that we have reviewed work on the functionalities of conversational search systems, we consider some fundamental properties of natural conversations. These properties could inform the design of conversational search systems to make them more cognitively compatible, and help assess what search functionalities are necessary to deliver the benefits offered by conversational interactions. Social and cognitive science provides rich insights into the general patterns of how people engage in conversations. Three basic principles are commonly recognized for governing these patterns [45]: *efficiency* [21, 50], *common ground* [35], and *recipient design* [9, 34, 50]. We advocate incorporating these principles into the design space of conversational search at each stage. Many of the functionalities discussed in the previous section can be seen as to serve these principles. In Table 1 we map them in this design space, and suggest additional functional goals that may be necessary to create a truly conversational experience. While each of these principles is backed

by a large volume of research, in the following we discuss high-level ideas and implications for designing conversational search.

Table 1 A design space of functional goals for conversational search

	Efficiency	Common ground	Recipient design
Query formulation	<ul style="list-style-type: none"> • Extract system input from natural language • Optimize elicitation dialogue flow • Maintain context and user model for inference 	<ul style="list-style-type: none"> • Clarification to bridge user input and system model • Incrementally suggest examples and elicit feedback 	<ul style="list-style-type: none"> • Support different querying behaviors • Tailor elicitation dialogue flow
Results exploration	<ul style="list-style-type: none"> • Minimize complexity • Start from summarization or low granularity, with navigational paths for details 	<ul style="list-style-type: none"> • Make capabilities and navigational paths discoverable • Support follow-up inquiries (paraphrasing, definition, examples, elaboration) • Suggest alternate information 	<ul style="list-style-type: none"> • Tailor search algorithms • Tailor presentation and interaction styles
Query reformulation and repair	<ul style="list-style-type: none"> • Maximize system initiative for repair • Support incremental repair 	<ul style="list-style-type: none"> • Signal breakdown • Make system model and status transparent • Suggest reformulation 	<ul style="list-style-type: none"> • Tailor repair assistance strategies

Efficiency: Also referred to as minimization, the efficiency principle guides speakers to achieve necessary informativeness with minimum effort, e.g., using as few turns and number of words as possible. To minimize user effort, this principle suggests the system to *maximize its initiative as allowed by the intelligence*, and *simplify the content as allowed by understandability*. System-initiated actions to support query formulation and repair, including extracting system input from user utterances, optimizing elicitation dialogues for user preferences, criteria or constraints, and contributing to repair, can be considered to serve the goal of minimizing user effort. Additionally, the system should in general aim to *support natural language expressions that may be minimal, implicit and incremental*, for example, by memorizing contexts and long-term user models to make inference about user’s information needs. For presenting information in a conversational form, this principle supports the idea of *starting from the lowest granularity of details that most targeted users can understand*, meanwhile providing rich navigational paths for details.

Common ground: This principle views conversation as a form of collective action to achieve mutual knowledge. Speakers constantly assess if there is good enough mutual understanding through evidence (e.g., explicit acknowledgment or relevant next turn), and if not, a grounding process (i.e., repair) will be initiated. When the dialogue partner is a machine, its model of understanding is significantly mismatched from the human speaker. Therefore, in applying the common ground principle for human-computer interactions, the general goal is to *make the system model more*

transparent, and actively *bridge between the user model and system model*. This principle is most relevant to functionalities for clarification and repair, and provides theoretical support for the three levels of system contributions we discussed for query repair: showing evidence (signaling potential misunderstanding or breakdown), making the system model (especially current understanding) transparent, and bridging mismatched user and system models (e.g., suggesting query reformulation). Equally important is to equip the system with capabilities to handle common types of user-initiated repair, including paraphrasing, providing definition, examples and elaboration [45]. The principle also highlights the importance of *system disclosure to help build users' expectation and mental models* of the search space at all stages, such as suggesting examples for feedback, recommending alternate information, and making the system capabilities and the navigational paths discoverable.

Recipient design: Speakers constantly tailor the ways they talk according to the particular recipient(s), based on their knowledge, social relations, personalities and so on. The principle of recipient design governs all aspects of human conversations, from choices of topics, to levels of details, to the organization of the conversation and linguistic styles. It is not surprising that incorporating awareness and knowledge of the user has long been a focus for research in dialogue systems and conversational agents [11], and many argue that demonstrating adaptiveness is necessary to achieve human-like conversational interactions. Work on conversational search, however, has not yet given much attention to this area. Of course, personalization and adaptation have their place in IR work [19, 28, 57], but mostly with a narrow focus on tailoring search algorithms based on users' interaction history or user profiles.

We argue that *recipient design should take a more central role for conversational interfaces*. First, besides *retrieval algorithms*, there are more properties of interactions to tailor for individual users to deliver a truly conversational experience. When presenting the search results, adapting *interaction styles* such as the level of details, language use, linguistic styles, and other social and communication designs for individual users could improve user engagement with agent systems [11, 56, 58, 65, 69]. During the process of query formulation and reformulation, many decision points in a dialogue flow, such as choices of elicitation methods, desired precision level, and repair strategies [2], could also account for individual differences.

Second, compared to traditional search systems where user interactions are limited to typing queries and selecting results, conversational interfaces allow users to give free-form natural-language input and thus enable much richer forms of user disclosure. Importantly, signals or user profiles obtained from the disclosure are the prerequisite for user modeling and system adaptation. In conversational interactions, the disclosure can not only be achieved through explicit inquiry, but also inference from users' conversational behaviors, just as how people engage in recipient design in everyday conversations based on subtle cues from the other speaker(s). These subtle cues may not only reveal who they are and what they prefer, but also how they feel about the interactions. In other words, conversational interfaces may enable new user modeling techniques for improving search experience based on feedback signals in users' conversational interactions that are beyond click-through patterns used by traditional IR systems. In our work developing conversational agents and

studying how users interact with these agents, we take great interest in identifying such user signals in conversational interactions to work towards the goal of system capabilities for recipient design. In the following section, we give an overview of our recent work to demonstrate some exciting opportunities in this area.

3 Recipient Design with a Conversational Search System

Starting in 2015, with a group of colleagues, we developed a conversational agent in IBM to answer employees' questions related to the work environment, such as "tell me about health benefits" or "how can I find IT help". The agent answers these questions using a hybrid model combining pattern-based retrieval of curated answers and output from a search engine for IBM internal web pages [12, 40]. The agent is called Cognitive Human Interface Personality (Chip). In summer 2016, Chip was deployed to 337 new hires who used Chip for 5-6 weeks. The deployment resulted in a chat log data set with more than 6,000 messages. The data set provides a valuable resource to study how users converse with a conversational agent *in the wild*.

We analyzed the data set with two goals in mind. First, we intended to provide an empirical account of the types of conversational interactions users have with an information-search agent. This complements prior work on conversational search in two ways. First, instead of conversing with a human intermediary or a wizard-of-oz system, we examined interactions with a real chatbot in a real-world setting. Second, besides information queries, we also paid attention to conversational interactions that perform communication or social functions, most of which are not covered by the existing models of information behaviors for conversational search [3, 49].

Our second goal was to explore recipient design, i.e. adapting the system, based on users' conversational behaviors, where we consider both tailoring its *search functionalities* and *interaction styles*. In this section we discuss insights from our work as an example to demonstrate the new opportunities that conversational interfaces offer for promoting user disclosure and hence enabling recipient design. The technical details of the system and the methodological details of the analysis are presented in our recent papers [12, 39, 40]. Before discussing the two areas of recipient design, we briefly describe the types of conversational interactions users had with Chip.

3.1 Conversational behaviors with a QA agent

Despite Chip being a QA agent for IBM internal information, we found that more than 35% of the chat logs were dedicated to non-queries. 85% of users sent at least one of these non-query messages. By performing content analysis, we categorized these non-query interactions into four general areas:

- **Opening and closing:** Instead of simply querying the agent, users showed anthropomorphizing behaviors by opening and closing the conversations as if

chatting with another human. More than 57% of users had at least once formally opened the interaction (e.g., “hi”, “hello”). 46% of users also had at least once closed the querying by acknowledging the agent’s answers (e.g. “ok”, “got it”), and 11.6% of users had at least once closed the conversations with farewell.

- **Agent ability checking:** There was a category of user questions concerned with capabilities of the system, by asking “what can you do” or “can you do [function]?”. These inquiries carried distinct meaning from other anthropomorphizing inquiries, as in serving the goal of reducing uncertainty about the system.
- **Feedback giving:** We found that users actively commented on the agent’s performance. This is interesting considering that it is a known challenge to obtain feedback in traditional search systems. During the deployment, we suggested participants to use “#fail” to give negative feedback if unsatisfied with Chip’s answers. 42.4% of users did it at least once. In addition, 11.9% users had at least once complimented Chip (e.g., “you are helpful”), and 21.1% made some forms of complaints (e.g., “you are stupid”).
- **Chit-chat:** Some users engaged in chit-chat with Chip, representing playful interactions by intentionally anthropomorphizing the agent [41]. Types of chit-chat included asking about the agent’s status (“what are you up to”) or traits (e.g. “what do you like?”), making off-topic requests, and talking about oneself.

Based on common patterns identified in Conversation Analysis, Moore et al. proposed a design framework for interaction patterns of conversational agents [45]. The framework differentiates between interactions that perform goal-oriented activities (inquire, respond, etc.), sequence management (e.g., repair), and conversation management. The main categories of conversation management interactions include opening, closing, capability checking, and disengaging (e.g., request to transfer to a human agent). The non-query conversational interactions we identified are generally consistent with these patterns of conversation management, with two additional areas—providing feedback and having playful chit-chat. Both can be considered unique to the setting of interactions with a personified conversational agent.

3.2 Recipient Design by Tailoring Search Functionalities

A longstanding theme in IR work is to adapt search functionalities to individual preferences based on feedback in the interaction history. For example, based on what a user liked or disliked, one can learn the user’s topical interest to adapt the ranking algorithm [19, 57]. Prior work also explored adapting information presentation or providing query assistance for those less satisfied users [64, 66]. Because it is costly and sometimes not feasible to obtain explicit user feedback, implicit feedback is often inferred from user behavior, such as click-through patterns or dwell time [17, 18, 28, 29]. However, the challenge is that these behavioral signals may be sparse and unreliable. In viewing conversational interactions users had with Chip, it was encouraging that users actively commented on its search performance (e.g., “#fail”).

It motivated us to further explore what feedback signals existed in their conversational interactions that can be leveraged for adapting search functionalities.

We adopted a data-driven approach by statistically modeling what features in conversational interactions predict a user’s self-reported satisfaction with Chip’s search performance, gathered by a survey. Details of the model are discussed in [40]. The results show that, after controlling for the system performance, users with lower *subjective* satisfaction tended to engage in conversational interactions in the categories of *feedback using #fail*, *agent ability checking*, *closing by farewell* and *off-topic requests*. These are the users that might have distinctive information needs and thus recipient design should target. A number of words or terms were also found to be associated with more positive user opinions, such as “tell me about”, “should I”, “what does”, “where is”, “who is”, “how to”, etc.

Based on these results, we summarize three areas in users’ conversational interaction to obtain feedback signals for adapting search functionalities: *conversational feedback*, *implicit complaints*, and *question structure keywords*:

- **Conversational feedback:** Conversational interfaces may encourage users to provide more feedback for the search performance. We found that using “#fail” was a strong indicator of dissatisfaction. In contrast, compliments such as “you are smart” or blunt complaints did not show significant association with user satisfaction, but instead might have been playful interactions. It highlights the complication in obtaining reliable feedback signals in conversational interactions, and the necessity to identify them by empirical analysis.
- **Implicit complaints:** A conversational interface enables users to freely “talk”, expressing feedback in ways that were not possible with traditional search systems. Statistical modeling allowed us to explore these less obvious signals. The result showed that the occurrences of *agent ability check* and *closing with farewell* predict user dissatisfaction. A closer look into the data revealed a pattern of users asking “what can you do” after encountering errors. Many have recognized that a critical drawback of conversational interface is its unclear affordance of capabilities [41, 52]. Agent ability-check can be considered signals of user struggling with such unclear affordance. Similarly, we observed users closing the conversation after errors, signaling frustration and refusal.
- **Question structure keywords:** It is reasonable to expect that users who are satisfied with the system performance are more likely to keep using it for information needs and less so for *off-topic requests*. Among the lexical features predicting user satisfaction, we saw a group of wh-words (what/where/who), which indicate typical questioning structures. With conversational interfaces, users tend to ask question in full sentences instead of using keyword-based queries [16, 23]. Therefore, one can possibly infer user satisfaction by tracking these structure keywords for questions.

This part of our work explored new opportunities offered by a conversational interface for user modeling to enable recipient design, i.e. adapting search results for individuals. We focused on informing ways to know *for whom* to adapt search functionalities. For addressing *how* the system should adapt, one can refer to the

large volume of IR work that tapped into user feedback signals to adapt search algorithms [5, 19, 57]. The general idea is that one could track user feedback signals for search results of different features (e.g., topics) to learn about user preferences. Future work could explore applying similar approaches to adapting features that are important for a conversational search setting, such as level of details in search results presentation. Another approach is to provide additional query assistance for less satisfied users, such as tailoring the dialogue flows for query elicitation and repair assistance strategies as suggested in Table 1.

3.3 Recipient Design by Tailoring Interaction Styles

In conversations, recipient design is reflected in more aspects than tailoring information content—in addition to *what* to talk about, one also constantly tailors *how* to talk. Therefore, besides search functionalities, we also attempted at adapting interaction styles for individual users. Interaction style is unarguably an enormous design space with many dimensions [56, 58, 65, 69]. We explored one dimension of core interest in studies of human-agent interaction—individuals’ orientation to view an agent as a sociable versus utilitarian tool. Prior work suggests that people tend to have different mental models interacting with a conversational agent [36, 37]. Those who see an agent as a sociable tool are inclined to engage in human-like interactions and exhibit relational behaviors such as chit-chat and politeness. Recent work also uses the term “playfulness” to refer to intentional engagement of anthropomorphizing behaviors with an agent system [41]. In contrast, those with a utilitarian orientation may see in an agent nothing more than an information search tool. Our hypothesis is that this orientation could govern a user’s preference for interaction styles that are more social versus ones that resemble traditional search systems.

To validate such a hypothesis and explore how to provide recipient design for social versus utilitarian interaction styles, we started with a qualitative study interviewing the users of our pilot deployment of Chip [39]. We developed a self-reported scale to measure one’s *social agent orientation* by asking if one enjoys conversational interactions and chit-chat with an agent. Then we contrasted user preferences for those on different sides of the scale. In general, we found that those with high social orientation *desire properties of natural conversations*, such as abilities to handle multi-turn conversation and tailoring the level of details for individual needs. They also prefer agents with *rich personality designs*. In contrast, those with utilitarian orientation repeatedly *favor common features in traditional search systems*, such as handling query-like input and providing a ranked list of answers. They also desire more *transparency of the information source*, and consider *human-like features to be unnecessary* in the conversation content and visual design.

The above results suggested *how* to tailor interaction styles for users with social versus utilitarian orientation. We also leveraged a statistical modeling approach to explore signals in conversational interactions to infer the orientation of an individual user, i.e. *for whom* to adapt the interaction style. We found the following categories

of conversational interactions to predict more social orientation: *chitchat about the agent's status* and *agent's traits*, *chitchat talking about oneself*, and *giving compliments*. We also found a number of words or terms associated with more social orientation: “how do you”, “are you”, “do you know”, “search”, “information”, etc.

We examine these results and summarize three areas to obtain signals for recipient design of social versus utilitarian interaction styles: *playful chitchat*, *agent-oriented conversations* and *casual testing*.

- **Playful chitchat:** Three categories of conversational interactions were strong signals of users with high social orientation—chitchat asking about the agent’s traits, status, and talking about oneself. They confirm that chit-chat carries explicit anthropomorphizing intentions. It is notable that conversation-management actions such as opening and acknowledging did not show significant association. It suggests that, in the context of a text-based QA agent, they may be more of habitual behaviors with the chat interface instead of consciously anthropomorphizing the agent.
- **Agent-orientated conversations:** An evident pattern in lexical features signaling social orientation is the frequent occurrence of second-person pronouns. This agent-oriented interest is consistent with the tendency to anthropomorphize the agent. This suggests that a simple way to identify socially-oriented users could be monitoring the usage of second-person pronouns.
- **Casual testing:** The lexical features predicting social orientation suggest less formality but more casual asking, such as “do you know” or “tell me”. We also found the words “information” and “search” to be strong signals. A close examination of the conversations revealed a pattern of repeatedly asking Chip to retrieve different kinds of information (e.g., “*search information about my manager*”). These behaviors suggest the less utilitarian-oriented users, who may see Chip differently from a traditional search system, exhibited curious behaviors by casually testing its intelligence.

With these studies, we aim to inform recipient design of conversational search with more social versus utilitarian interaction styles. For example, based on the above signals, one can distinguish users with social or utilitarian orientation, interacting with natural conversations and rich social designs for the former group, while presenting information like traditional search system for the latter. Besides recipient design for search results presentation, one may also improve the search task by anticipating differences in users’ querying behaviors. For example, it is likely that a socially oriented user would express information needs in natural language, while a utilitarian oriented user would use keywords based queries, and thus different processing techniques should be applied. Our work demonstrates that recipient design should take a more central role for conversational interfaces because the rich design issues involved, for not only improving search algorithms, but also interaction styles, which could be concerned with information presentation, conversation organization, or linguistic styles.

4 Summary

Conversational interfaces are an emerging area of research for search systems. It is important to recognize that human conversation is a metaphor for this type of user interface. The benefit of an interface metaphor is to give users instantaneous knowledge on how to interact with the system in a familiar way, while leaving the interface to bridge these familiar actions and the underlying computational models. To take full advantage of such a metaphor, this chapter aims to provide perspectives on properties of natural conversations that fit and benefit information search tasks, in order to identify functional goals of conversational search systems. By reviewing relevant work, we start from two ends: system actions to extend models of traditional search systems to conversational search, with regard to query formulation, search results exploration, and query repair; and the fundamental properties of natural conversations including efficiency, common ground and recipient design. The two threads converge at a design space for the functional goals for conversational search systems. We identify a gap and an area of opportunity to put user modeling and adaptation in a more central place for conversational interfaces, and discuss insights from our own work on making search systems conversational and adaptive.

References

1. Arguello J, Choi B, Capra R. (2017). Factors Affecting Users' Information Requests. In SIGIR 1st International Workshop on Conversational Approaches to Information Retrieval (CAIR'17) (Vol. 4).
2. Ashktorab Z, Jain M, Liao QV, Weisz, J (2019). Resilient Chatbots: Repair Strategy Preferences for Conversational Breakdown. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI'19). ACM, New York, NY, USA
3. Azzopardi L, Dubiel M, Halvey M, Dalton J. (2018). Conceptualizing agent-human interactions during the conversational search process. In The Second International Workshop on Conversational Approaches to Information Retrieval.
4. Balchandran R, Rachevsky L, Sansone L, Sicconi R. (2009, September). Dialog system for mixed initiative one-turn address entry and error recovery. In Proceedings of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue (pp. 152-155). Association for Computational Linguistics.
5. Baeza-Yates R, Ribeiro BAN. (2011). Modern information retrieval. New York: ACM Press; Harlow, England: Addison-Wesley,.
6. Bates, M. J. (2002). Toward an integrated model of information seeking and searching. *The New Review of Information Behaviour Research*, 3, 1-15.
7. Baeza-Yates R, Ribeiro BAN. (2011). Modern information retrieval. New York: ACM Press; Harlow, England: Addison-Wesley,.
8. Belkin NJ, Cool C, Stein A, Thiel U. (1995). Cases, scripts, and information-seeking strategies: On the design of interactive information retrieval systems. *Expert systems with applications*, 9(3), 379-395.
9. Bell A. (1984). Language style as audience design. *Language in society*, 13(2), 145-204.
10. Bobrow, D. G., Kaplan, R. M., Kay, M., Norman, D. A., Thompson, H., Winograd, T. (1977). GUS, a frame-driven dialog system. *Artificial intelligence*, 8(2), 155-173.

11. Cassell J, Sullivan J, Churchill E, Prevost S (Eds.). (2000). Embodied conversational agents. MIT press.
12. Chandar P, Khazaeni Y, Davis M, Muller M, Crasso M, Liao QV, ... Geyer W. (2017). Leveraging conversational systems to assist new hires during onboarding. In IFIP Conference on Human-Computer Interaction (pp. 381-391). Springer, Cham.
13. Chen L, Pu P. (2004). Survey of preference elicitation methods (No. EPFL-REPORT-52659).
14. Christakopoulou K, Radlinski F, Hofmann K. (2016). Towards conversational recommender systems. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 815-824). ACM.
15. Christakopoulou K, Beute, A, Li R, Jain S, Chi EH. (2018). Q&R: A two-stage approach toward interactive recommendation. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 139-148). ACM.
16. Crestani F, Du H. (2006). Written versus spoken queries: A qualitative and quantitative comparative analysis. *Journal of the American Society for Information Science and Technology*, 57(7), 881-890.
17. Feild HA., Allan J, Jones R. (2010). Predicting searcher frustration. In Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval (pp. 34-41). ACM.
18. Fox, S, Karnawat, K, Mydland, M, Dumais, S, White, T. (2005). Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)*, 23(2), 147-168.
19. Ghorab MR., Zhou D, O'connor A, Wade V (2013). Personalised information retrieval: survey and classification. *User Modeling and User-Adapted Interaction*, 23(4), 381-443.
20. Goodwin C, Heritage J. (1990). Conversation analysis. *Annual review of anthropology*, 19(1), 283-307.
21. Grice HP. (1975). Logic and conversation. In Cole, P.; Morgan, J. *Syntax and semantics*. 3: Speech acts. New York: Academic Press. 41-58.
22. Gunning D. (2017). Explainable artificial intelligence (xai). Defense Advanced Research Projects Agency (DARPA), nd Web
23. Guy I. (2018). The Characteristics of Voice Search: Comparing Spoken with Typed-in Mobile Web Search Queries. *ACM Transactions on Information Systems (TOIS)*, 36(3), 30.
24. Hassan AA, Gurunath KR, Ozertem U, Jones R. (2015). Characterizing and predicting voice query reformulation. In Proceedings of the 24th ACM International on Conference on Information and Knowledge Management (pp. 543-552). ACM.
25. Heritage J, Atkinson, JM. (1984). Structures of social action. *Studies in Conversation Analysis*.
26. Jiang J, Jeng W, He D. (2013). How do users respond to voice input errors?: lexical and phonetic query reformulation in voice search. In Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval (pp. 143-152). ACM.
27. Jiang J, Hassan AA, Jones R, Ozertem U, Zitouni I, Gurunath KR, Khan OZ. (2015). Automatic online evaluation of intelligent assistants. In Proceedings of the 24th International Conference on World Wide Web (pp. 506-516). International World Wide Web Conferences Steering Committee.
28. Kelly D, Teevan J. (2003). Implicit feedback for inferring user preference: a bibliography. In *Acm Sigir Forum* (Vol. 37, No. 2, pp. 18-28). ACM.
29. Kim Y, Hassan A, White RW, Zitouni I. (2014). Modeling dwell time to predict click-level satisfaction. In Proceedings of the 7th ACM international conference on Web search and data mining (pp. 193-202). ACM.
30. Mahmood T, Ricci F. (2009). Improving recommender systems with adaptive conversational strategies. In Proceedings of the 20th ACM conference on Hypertext and hypermedia (pp. 73-82). ACM.
31. McGinty L, Smyth B. (2006). Adaptive selection: An analysis of critiquing and preference-based feedback in conversational recommender systems. *International Journal of Electronic Commerce*, 11(2), 35-57.
32. McTear MF. (2004). Spoken dialogue technology: toward the conversational user interface. Springer Science and Business Media.

33. McTear M, Callejas Z, Griol D. (2016). *The conversational interface: Talking to smart devices*. Springer.
34. Clark HH, Murphy GL. (1982). Audience design in meaning and reference. In *Advances in psychology* (Vol. 9, pp. 287-299). North-Holland.
35. Clark, HH, Brennan, SE. (1991). Grounding in communication. *Perspectives on socially shared cognition*, 13(1991), 127-149.
36. Lee MK, Kiesler S, Forlizzi J. (2010). Receptionist or information kiosk: how do people talk with a robot?. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work* (pp. 31-40). ACM.
37. Lee N, Shin H, Sundar, SS. (2011). Utilitarian vs. hedonic robots: role of parasocial tendency and anthropomorphism in shaping user attitudes. In *Proceedings of the 6th international conference on Human-robot interaction* (pp. 183-184). ACM.
38. Li R, Kahou SE, Schulz H, Michalski V, Charlin L, Pal C. (2018). Towards Deep Conversational Recommendations. In *Advances in Neural Information Processing Systems* (pp. 9748-9758).
39. Liao QV, Davis M, Geyer W, Muller M, Shami NS. (2016). What can you do?: Studying social-agent orientation and agent proactive interactions with an agent for employees. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (pp. 264-275). ACM.
40. Liao QV, Hussain MM, Chandar P, Davis M, Crasso M, Wang D, ... Geyer W. (2018). All Work and no Play? Conversations with a Question-and-Answer Chatbot in the Wild. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI'18)*. ACM, New York, NY, USA (Vol. 13).
41. Luger E, Sellen A. (2016). Like having a really bad PA: the gulf between user expectation and experience of conversational agents. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 5286-5297). ACM.
42. Marchionini, G. (1997). *Information seeking in electronic environments* (No. 9). Cambridge university press.
43. McDuff D, Thomas P, Czerwinski M, Craswell N. (2017). Multimodal analysis of vocal collaborative search: a public corpus and results. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction* (pp. 456-463). ACM.
44. Miller T. (2018). *Explanation in artificial intelligence: Insights from the social sciences*. Artificial Intelligence.
45. Moore RJ, Arar R. (2019). *Conversational UX Design: A Practitioner's Guide to the Natural Conversation Framework*. Morgan & Claypool.
46. Myers C, Furqan A, Nebolsky J, Caro K, Zhu J. (2018). Patterns for How Users Overcome Obstacles in Voice User Interfaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (p. 6). ACM.
47. Paek T, Horvitz E. (2000). Conversation as action under uncertainty. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence* (pp. 455-464). Morgan Kaufmann Publishers Inc..
48. Porcheron M, Fischer JE, Reeves S, Sharples S. (2018). Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (p. 640). ACM.
49. Radlinski F, Craswell N. (2017). A theoretical framework for conversational search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval* (pp. 117-126). ACM.
50. Sacks H, Schegloff, E A. (1979). Two preferences in the organization of reference to persons in conversation and their interaction. In G. Psathas (ed) *Everyday language: Studies in ethnomethodology*, New York.
51. Schegloff EA. (2007). *Sequence organization in interaction: Volume 1: A primer in conversation analysis* (Vol. 1). Cambridge University Press.
52. Shneiderman B, Maes P (1997). Direct manipulation vs. interface agents. *interactions*, 4(6), 42-61.
53. Shokouhi M, Jones R, Ozertem U, Raghunathan K, Diaz F. (2014). Mobile query reformulations. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval* (pp. 1011-1014). ACM.

54. Sitter S, Stein A. (1992). Modeling the illocutionary aspects of information-seeking dialogues. *Information Processing & Management*, 28(2), 165-180.
55. Sun Y, Zhang Y. (2018) Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM.
56. Szafr D, Mutlu B. (2012). Pay attention!: designing adaptive agents that monitor and improve user engagement. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 11-20). ACM.
57. Teevan J, Dumais ST, Horvitz E. (2005). Personalizing search via automated analysis of interests and activities. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval* (pp. 449-456). ACM.
58. Thomas P, Czerwinski M, McDuff D, Craswell N, Mark G. (2018). Style and Alignment in Information-Seeking Conversation. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (pp. 42-51). ACM.
59. Thompson CA, Goker MH, Langley P. (2004). A personalized system for conversational recommendations. *Journal of Artificial Intelligence Research*, 21, 393-428.
60. Trippas JR., Spina D, Cavedon L, Joho H, Sanderson M. (2018). Informing the Design of Spoken Conversational Search: Perspective Paper. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (pp. 32-41). ACM.
61. Walker M, Aberdeen J, Boland J, Bratt E, Garofolo J, Hirschman L, ... Pellom B. (2001). DARPA Communicator dialog travel planning systems: The June 2000 data collection. In *Seventh European Conference on Speech Communication and Technology*.
62. Wang WY, Finkelstein S, Ogan A, Black AW, Cassell J. (2012 July). Love ya, jerkface: using sparse log-linear models to build positive (and impolite) relationships with teens. In *Proceedings of the 13th annual meeting of the special interest group on discourse and dialogue* (pp. 20-29). Association for Computational Linguistics.
63. Wilson, T. D. (2000). Human information behavior. *Informing science*, 3(2), 49-56.
64. Song Y, He, L W. (2010). Optimal rare query suggestion with implicit user feedback. In *Proceedings of the 19th international conference on World wide web* (pp. 901-910). ACM.
65. Xiao J, Stasko J, Catrambone R. (2007). The role of choice and customization on users' interaction with embodied conversational agents: effects on perception and performance. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 1293-1302). ACM.
66. Zhai C, Lafferty J. (2006). A risk minimization framework for information retrieval. *Information Processing & Management*, 42(1), 31-55.
67. Zhang Y, Chen X, Ai Q, Yang L, Croft WB. (2018). Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (pp. 177-186). ACM.
68. Zhang Y, Liao QV, Srivastava, B. (2018, March). Towards an Optimal Dialog Strategy for Information Retrieval Using Both Open-and Close-ended Questions. In *23rd International Conference on Intelligent User Interfaces* (pp. 365-369). ACM.
69. Zhao R, Sinha T, Black A, Cassell J. (2016). Automatic recognition of conversational strategies in the service of a socially-aware dialog system. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue* (pp. 381-392).