

Towards System-Initiative Conversational Information Seeking

Somin Wadhwa, Hamed Zamani

Center for Intelligent Information Retrieval
University of Massachusetts Amherst
Amherst, MA, United States

Abstract

Presently, most conversational information seeking systems function in a passive manner, i.e., user-initiative engagement. Through this work, we aim to discuss the importance of developing conversational information seeking systems capable of system-initiative interactions. We further discuss various aspects of such interactions in CIS systems and introduce a taxonomy of system-initiative interactions based on three orthogonal dimensions: initiation moment (*when to initiate a conversation*), initiation purpose (*why to initiate a conversation*), and initiation means (*how to initiate a conversation*). This taxonomy enables us to propose a generic pipeline for system-initiative conversations, consisting of three major steps associated with the three dimensions highlighted in the taxonomy. We further delineate the technical and evaluation challenges that the design and implementation of each component may encounter, and provide possible solutions. We finally point out potential broader impacts of system-initiative interactions in CIS systems.

Keywords

Conversational search, conversational information seeking, mixed-initiative conversations, conversational recommendation

1. Introduction

The rapid growth in speech and small screen interfaces has significantly influenced the way users interact with intelligent systems to satisfy their information needs. The growing interest in personal digital assistants demonstrates the willingness of users to employ conversational interactions. This has motivated the information retrieval community, both academic researchers and industry practitioners, to focus on conversational information seeking (CIS) as a major emerging research area.¹ It has been also recognized as one of the strategic directions of the community in the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018) [1].² However, current models and technology provide limited support to conversational understanding and various types of interactions. Recent research has made substantial progress in a number of tasks associated with conversational information seeking [2, 3, 4, 5], however, each with various simplifying assumptions on system abilities and user behavior that may not hold in a real-world CIS system [6, 7]. For instance, mixed-initiative interactions have been largely

ignored in most recent work in the area of conversational information seeking. This is while mixed-initiative intelligent systems are believed to ultimately revolutionize the world of computing [7], and CIS systems provide an appropriate platform for supporting mixed-initiative interactions.

Recently, some form of such interactions have been studied in the context of asking for clarification [8, 9, 10] or preference elicitation [11, 12]. Developing fully mixed-initiative conversational systems requires support for *system-initiative* (or agent-initiative) interactions, where the CIS system initiates a conversation with the user(s). However, system-initiative interactions have been overlooked in the CIS literature. In this paper, we focus on this topic and discuss its importance for IR research and industry. We believe that real-life intelligent assistants can substantially benefit from supporting system-initiative interactions and this direction involves a large number of unsolved and non-trivial open questions that are worthy of research. To better demonstrate different aspects of the problem, we compile a taxonomy of system-initiative interactions, based on three dimensions: (1) initiation moment: *when* to initiate a conversation, (2) initiation purpose: *why* to initiate a conversation, and (3) initiation means: *how* to initiate a conversation. We believe that system-initiative interactions can be categorized as either instant initiation or opportune moment initiation interactions. We provide example scenarios for each of these categories in Section 2.

The introduced taxonomy enables us to propose a generic pipeline for system-initiative interactions in CIS systems. The pipeline introduced in Section 3 consists

DESIREs 2021 – 2nd International Conference on Design of Experimental Search & Information REtrieval Systems, September 15–18, 2021, Padua, Italy

✉ sominwadhwa@cs.umass.edu (S. Wadhwa);
zamani@cs.umass.edu (H. Zamani)

© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

¹In this paper, we use CIS to refer to all conversational information seeking and access systems, including conversational search, recommendation, and question answering.

²<https://sites.google.com/view/swirl3/>

of three major steps, that are aligned with the three dimensions in our taxonomy. We further review technical challenges in both modeling and evaluating each of these steps in addition to discussing potential approaches for end-to-end evaluation of system-initiative CIS systems. We also highlight the dangers of using system-initiative interactions in CIS systems if not designed carefully. We finally briefly introduce the broader impact of this research direction. We believe this paper, despite being sometimes abstract or hypothetical, sheds light on some aspect of developing and evaluating system-initiative conversational information seeking systems.

2. A Taxonomy of System-Initiative CIS Interactions

In this section, we review different interactions that may be taken by a CIS system to initiate a conversation. We study these interactions with respect to the following three orthogonal dimensions:

- initiation moment: *when* to initiate a conversation?
- initiation purpose: *why* to initiate a conversation?
- initiation means: *how* to initiate the conversation?

We believe that any CIS system should be able to answer all the above questions in order to make system-initiative interactions. In the rest of this section, we explain these dimensions. This paper also proposes a pipeline for system-initiative interactions in CIS systems, which is inspired by these three dimensions introduced in the taxonomy.

2.1. Dimension I: Initiation Moment

Given the first dimension, i.e., when to initiate a conversation, we partition system-initiated conversational interactions into two categories:

- **Instant initiation:** defined as instant initiation of a conversation is by a conversational information seeking system mostly based on the user's current situation.
- **Opportune moment initiation (OMI):** defined as initiation of a conversation that can be postponed to an opportune moment that is decided by the conversational information seeking system.

In other words, the first category contains the interactions that should be initiated instantly and are not appropriate in other contexts. The second category, on

the other hand, contains the interactions that can be initiated at a later time decided by the system.³ Therefore, the interaction time in instant initiation is derived by the user's situational context, e.g., user's location, time, mood, and activity, or the urgency of the interactions (e.g., health and safety related interactions), while in OMI, this is the CIS system that decides the interaction time.

2.2. Dimension II: Initiation Purpose

Conversation initiation may be triggered by availability of a *new data* that may be of interest to user, by the *current situation* of user such as time and location, or by *modifications to the CIS system*. The latter may happen for example if a new deployment of the CIS models leads to an understanding that the system provided false information to a sensitive topic in the past interactions and now wants to initiate a conversation to correct its past mistake. Given these three triggering reasons, we identify five main purposes for initiating a conversation in a CIS system. They include information filtering, recommendation, following up a past conversation, contributing to multi-party conversation, and feedback request. Note that this paper only focuses on information seeking conversations, therefore there exist some non information seeking initiation purposes that are not covered in this section.

In the following, we describe each of the identified initiation purposes. For each initiation purpose presented below, we provide instant initiation and opportune moment initiation example use-cases in Table 1.

Filtering streaming information Information filtering systems aim for delivering information to the user from a *stream of information contents* based on the user's preferences. Belkin and Croft [13] identified information retrieval and information filtering as two sides of the same coin, because of their fundamental similarities in representing unstructured or semi-structured documents and computing their relevance to the user's (short- or long-term) information needs. A few years later, Robertson and Hull [14] organized the TREC Filtering Tracks to promote the field and provide resources for fostering research in the filtering tasks. Conversational information seeking systems may initiate a conversation with the goal of information filtering. For instance, introducing the breaking news headlines based on the user's preferences is considered as an information filtering task that may have applications in system-initiative CIS systems.

Recommendation Recommender systems are often considered as a subcategory of information filtering sys-

³OMI interactions can also be triggered by the user at a convenient time.

Table 1

Examples for various initiation purposes (rows) based on initiation moments (columns).

	Instant Initiation	Opportune Moment Initiation
Filtering streaming information based on user profile	Health and safety related information is often time-sensitive. For instance, attacks or events that may lead to a safety risk or hazard for the user should be instantly mentioned by a CIS system that is watching these streaming information sources.	News agencies are constantly publishing new content on their website. Users, on the other hand, have different preferences and tastes in the news topics and sources. A system-initiative CIS system may initiate a conversation, based on the opportune moment initiation scheme, to inform the user based on their preferences.
Recommendation	Many users create and maintain to-do lists for their daily activities. A few recent recommender systems have been developed to re-rank and recommend the next to-do item. Some of the items in a to-do list can be time-sensitive and a CIS system can instantly initiate a conversation to notify the user that the deadline for doing one of the yet-to-be-done tasks in the to-do list is approaching, otherwise the user will not be able to complete the task.	Active engagement through CIS can also occur in broad opportune moments like the pre-holidays. People are often known to exchange gifts during some special occasions and holidays and a CIS could play an active role in offering gifting recommendations to the user. Such an active engagement would be time-sensitive, and in addition to user-preferences for gift recommendations, a window-of-initiation would be equally as relevant.
Following up a past user-system conversation	Any modification to the system's response for a health or safety related question of the user which was asked in the past may need a prompt conversation initiation. For instance, if the user asks about the number of daily COVID-19 cases in an institute, and the system responds with zero, it may need to instantly initiate a conversation upon discovering a new case in the day. (note that many examples in this category also involve filtering of streaming information, however such filtering should happen with respect to the past user-system interactions, which is different from the first row in this table.)	CIS systems are not by any means perfect and they make mistakes in responding to user's requests. Based on new information or new models deployed in the system, a CIS system may initiate a conversation at an opportune moment to accept and correct its mistakes that was made in the past.
Contributing to a multi-party human conversation	While it is largely unexplored in the literature, one possible use-case of a system-initiative CIS engagement in a human-human interaction could be that of monitoring the factual accuracy of the underlying content exchanged in human conversations (if and where necessary). The CIS system may engage in retrieval-based fact-checking and initiate a conversation to contribute to the ongoing human conversation by providing the fact-checking results and details.	Similar to the previous case with a focus on the monitored past human conversations (i.e., following up a past human conversation).
Feedback request	Asking for a location- and time-specific feedback may need to happen promptly. For example, while a user is driving and passing by a specific location, a CIS system may initiate a conversation for feedback request by asking about a car accident in that location.	An example of an opportune moment feedback request is that of e-commerce shopping. Under the current popular systems, users often indiscriminately required to provide reviews of products right after they purchase them or after a pre-defined period of time. Factoring-in the category of products along with user meta-data could enhance a CIS's ability to gauge what moments would be most opportune in terms of engaging an active conversation about seeking product feedback.

tems, however, we intentionally separate these two in this paper to highlight their differences and important applications in system-initiative CIS systems. Unlike information filtering tasks that deal with a stream of data, in this paper, recommendation tasks refer to recommending entities or information from an existing data source. For instance, recommending a restaurant based on the user’s location and preferences can be considered as a recommendation task but does not fit well within the definition of information filtering tasks provided above. CIS systems may initiate a conversation to make a recommendation to the user.

Following up a past conversation A CIS system may follow up a past conversation for many different reasons, such as providing new information that was not available at the time of past conversation, correcting a mistake that was made by the system in a past conversation, and continuing a conversation that was interrupted and left incomplete. System-initiation enables CIS systems to follow up past conversations to better serve their ultimate information seeking and access purpose.

Contributing to a multi-party human conversation Existing conversational information seeking systems are mainly designed for user-system interactions. However, CIS systems can contribute to multi-party human conversations, such as collaborative conversations. For instance, based on a conversation between two people, a CIS system that is permitted to monitor the conversation may contribute to the topic of the discussion, e.g., by fact-checking the claims made in the conversation and taking an initiative if a false claim is made by one party.

Feedback request Feedback requests are not directly related to information seeking, however, user’s feedback, such as product reviews, plays a key role in development of several information seeking systems. On the other hand, users often forget or refuse to provide a feedback. In some cases, a CIS system may initiate a conversation with the goal of collecting feedback about the user’s experiences. Such conversation may convince the users to provide feedback in cases where they normally do not.

2.3. Dimension III: Initiation Means

How to initiate a conversation shapes the third dimension in our conversation initiation taxonomy. In case of multi-device setting, the system should decide which device should be used to initiate a conversation. Or in case of multi-modal setting, the system should decide which interaction channel (e.g., visual through a screen or aural through the speaker) or processing modality (e.g., verbal through text or non-verbal through an image) should be used for initiating a conversation. One

Table 2
Notation descriptions.

symbol	description
u	the user
p_u^t	the user profile and situational context associated with u at timestamp t
c_u^t	all the conversational interactions of u with the CIS system up to timestamp t
C^t	The collection of all information items available at timestamp t (e.g., from the web)
i	a system initiation instance object
\mathcal{D}	a collection of system initiation instance objects

can imagine a system that can ask for a permission to initiate a conversation, for example via a light vibration.

3. A Pipeline for Conversation Initiation in CIS

As mentioned in the last section, information seeking conversations can be initiated by new information, by the situational user context, and by new model deployment. In this section, we present a general high-level pipeline for initiating a conversation in CIS systems. Due to the complexity of developing and evaluating the pipeline for system-initiative interactions, we additionally provide a formal definition of each step. This formalization enables us to easily discuss evaluation methodologies for each component in Section 6. It also helps future work to see these steps in isolation. The pipeline is depicted in Figure 1. It consists of the following steps that use the notation introduced in Table 2.

Step I: Producing system initiation instances In the first step, system initiation instances are produced by the processes described in Section 2.2, such as recommendation and contributing to a multi-party conversation. They are shown as *initiation purposes* in Figure 1. These processes monitor the environment and produce instances that can lead to system initiation by observing new filtered information or recommendation based on the user’s context (see Section 2 for more detail about these processes). The produced conversation initiation instances will be added to the instances collection (or database). Note that a system initiation instance is a data object that contain all the information required for initiating a conversation, including the initiation purpose, the data, context, or reason that led to the production of the instance, the initiation features and content, etc. This step can be formalized as a function of p_u^t , c_u^t , and C^t that produces one or more system initiation instances,




Figure 1: A generic pipeline for conversation initiation in CIS systems. Each initiation instance is a data object containing all information required for initiating a conversation, including the initiation purpose, content, and context.

i.e., $\phi(p_u^t, c_u^t, C^t)$.

Step II: Selecting an instance for conversation initiation In the next step, the initiator component (see Figure 1) selects one of the entries in the instances collection \mathcal{D} for initiating a conversation. Although some conversations need to be initiated promptly (i.e., instant initiation), in this pipeline all instances are inserted into \mathcal{D} and this is the job of the initiator to promptly identify instant initiation requests. In more details, the initiator component is constantly monitoring all entries in \mathcal{C} and based on the user’s situational context decides what instance should be selected at each timestamp for conversation initiation. This step can be formalized as $\psi(i, p_u^t) = \Pr(\text{initiation} = 1 | i, p_u^t)$ where i is a system initiation instance in \mathcal{D} and “initiation” is a binary hidden variable representing the event of initiating the conversation. Note that although everything mentioned in this paper is about system-initiative conversations, note that the initiator can be also triggered by the user (for instance, the user may say “I’m board, tell me something”).

Step III: Conversation generation Once the initiator component selects one of the instances from \mathcal{D} , a natural language utterance will be generated by the conversation generation component and it will be presented to the user based on an appropriate device and interaction modality (in case of multi-device or multi-modal settings). Therefore, this step can be formally defined as a function that generates a conversation based on a given instance i and presented to the user based on the

user profile and situational context, i.e., $\gamma(i, p_u^t)$.

4. User Response to System-Initiative Conversations

While users are free to respond in any form they may see fit, for a substantive functioning of the system we propose a certain categorization of responses based on how they are processed by the system:

- **Null action:** User provides no response to the initiated conversation by the CIS system. Note that null action should not necessarily be interpreted as a negative feedback, since the user may find some initiation useful, while they are not interested in further engagement.
- **Interruption or negation:** User provides a response consistent with the interpretation of shutting down any further engagement by the CIS system. Such response can be safely assumed as a negative feedback.
- **Relevant response:** User responds to the initiated conversation by a relevant answer. This is often expected to happen when the initiated conversation involves a question or asks for feedback.
- **Postpone:** User responds to the initiated conversation and asks the system to remind them at a later time.
- **Critique or clarification-seeking response:** The kind of responses here would include users further engaging

in a back-and-forth conversation with agent about either seeking further information or critiquing existing engagement. One key technical challenge here that we talk about in the next section would be the processing of the user response in order to inculcate it to make the system better.

- Follow up: User responds with a follow up response to get further information or perform actions related to the initiated conversation.
- Topic drift: User responds but changes the topic of the initiated conversation.

Given the current status of text classification models and the complexity of the task, it is possible to achieve an acceptable classification accuracy in classifying user responses to the above categories. They can be further used for training or evaluating the conversation initiation process. For instance, interruption or negation may be considered as negative feedback. Such feedback can be used to modify the models deployed for each of the three steps (ϕ , ψ , and γ) in the pipeline (see Section 3). On the other hand, receiving a relevant response may be considered as a positive feedback for the system.

5. Technical Challenges

In this section, we hypothesize certain key technical challenges in implementing the pipeline described in Section 3.

5.1. Producing System-Initiative Instances

The first step in the system-initiation pipeline is to identify reasons for initiating a conversation and generate a system-initiative instance. As described in the last section, system-initiative instances are data objects that contain all the information about a system-initiative conversation, such as the purpose, content, and context. This step can be cast to implementing each of the five initiation purpose components discussed in Section 2.2. In other words, one needs to implement the function $\phi(p_u^t, c_u^t, C^T)$ with a focus on each initiation purpose. This has roots in various IR tasks, such as filtering and recommendation. However, some of the initiation purposes are relatively unstudied in the literature, such as following up a past conversation or contributing to a multi-party conversation. Even feedback request in the form of active conversation is underexplored. Therefore, one of the major technical challenges in producing system-initiative instances is to develop models that can identify the reasons for conversation initiation when the

goal is either filtering of streaming information, recommendation, conversation follow-up, contributing to a multi-party conversation, or feedback request.

5.2. Developing an Initiator Model

The second step in the provided pipeline is selecting a system-initiative instance from the instance collection \mathcal{D} by an initiator component (see Figure 1). This is indeed equivalent with implementing the function $\psi(i, p_u^t)$. We believe that the most challenging part of implementing such component is our lack of knowledge on what is generally the right moment for initiating a conversation. Therefore, we believe that future research should focus on conducting user studies in the wild to explore what are the right time to initiate a conversation. Some weak supervision signals can be mined from user interactions with the current conversational systems, even if they do not support system-initiative interactions. For example, the times when the user initiate an unimportant conversation (due to being board for example) can provide a weak (noisy) signal as a potentially good time to initiate a conversation and thus machine learning based models can be trained based on the situational context and the user profile to predict such moments. Of course, a nice property of interactive systems that log the user interactions is to iteratively improve the system ability to accurately predict such moments based on the feedback received from the user (see Section 4 for various types or user responses to system-initiative interactions).

5.3. Generating System-Initiative Utterances

The third and the final step in our pipeline (Section 3) is to generate a conversation based on a system-initiative instance and present it to the user, equivalent to implementing the function $\gamma(i, p_u^t)$. We believe that many techniques developed in the dialogue systems and text generation research can be used for implementing this component. Each instance i is a structured data object, therefore, neural models for unstructured text generation from structured data, e.g., tables, can be potentially adopted. Since the users mostly do not expect system-initiative utterances, an interesting technical challenge here would be providing some context in the generated utterance to make sure that the user understands why such conversation being initiated. This context may refer to a previous interaction of the user with the system, a past experience of the user, or an explanation on the reason that led to the generation of such system-initiative conversation.

6. Evaluating System-Initiative CIS Systems

Evaluation is one of the most challenging aspect of system-initiative CIS systems. IR research has a long history of collection creation for various information seeking tasks, however, they are mostly created based on a set of pre-defined information needs (e.g., most TREC tracks) or a set of observations (e.g., clickthrough data). Such evaluation methodologies do not easily extend to an active interaction scenarios, such as system-initiation in conversation.

Although evaluating system-initiative CIS systems is yet to be explored in the literature, in this section, we detail our perspective on potential evaluation methodologies that can be pursued.

6.1. Evaluating system initiation instances (ϕ)

As pointed out in Section 3, the first step towards initiating a conversation is to produce system initiation instanced formalized using a function $\phi(p_u^t, c_u^t, C^T)$. The initiation command should include all information about the nature of conversation initiation. To evaluate this component, we should provide all the required information at the timestamp t to the system as input and evaluate the produced initiation command. The required information (as depicted in Figure 1) includes past user-system interactions, user profile, user situational context, and a stream of new information content. The model should produce an initiation command or NULL, meaning that no initiation is needed. Both precision and recall-oriented metrics should be used to evaluate the model's performance. In fact, the produced initiation instances should be of high quality (all initiations should be relevant to the user) with high coverage (all required initiations should be produced by the model).

Based on this evaluation methodology, a reusable collection can be created. The data collection can be either sampled from a real user's interaction history (a realistic setting, but requires access to real user-system conversational interactions), or constructed based on information seeking interactions between two or more people. The latter can be done in a lab study using a wizard of oz setting, similar to [15].

6.2. Evaluating Initiation Moments (ψ)

To evaluate the initiator component in the proposed pipeline (see Figure 1), one can cast the problem to a binary classification task. In more detail, we can formalize the task as predicting whether to initiate the conversation or not given an initiation instance and a set of

information related to the user's profile and context (e.g., time and location), formalized as $\psi(i, p_u^t)$ in Section 3. This approach assumes that the importance of initiation moment is binary (good or bad). However, this is not the case. Two situations may be bad for initiating a conversation, but one may be the worst. Therefore, we believe that this task should be evaluated as a ranking task: re-rank the list of situational context information associated with a user given for an initiation command. This setting allows us to have multi-level (or graded) labels for each situational context and use relevant metrics (e.g., similar to NDCG [16]) to evaluate the quality of the system in identifying the right situation (moment) to initiate the conversation.

6.3. Evaluating the Content of Initiated Conversations (γ)

After identifying conversation initiation commands, the system needs to produce a natural language sentence or utterance (in most cases) and select the initiation means if needed for initiating the conversation. In Section 3, we formalize this as $\gamma(i, p_u^t)$. To evaluate this ability of the system, we can assume that the initiation commands are accurately produced and are complete (i.e., a hypothetically ideal system with perfect precision and recall). Based on this assumption, the focus of this evaluation step would be to generate a conversation utterance based on a given initiation instance i . The generated utterance should contain all required information in addition to being precise and fluent. A number of ground truth reference utterances can be generated through manual annotation and popular text generation metrics such as BLEU [17], ROUGE [18], and BERTScore [19] may be used to evaluate the model. As discussed in [20], despite the popularity of these metrics, they do not necessarily reflect the quality of the produced dialogue, and ultimately human annotation of the model's outputs is desired.

6.4. End-to-End Evaluation of System-Initiated Conversations

The last three subsections discuss component-level evaluation of system-initiative CIS systems. As mentioned above, each is based on some simplifying assumptions of other components of the system, which is unrealistic. Therefore, an end-to-end evaluation of system-initiated conversations should be explored. To do so, both offline and online evaluation strategies can be adopted. For offline evaluation, each instance would include all the required information for the system at a timestamp t as input, including past user-system interactions, user profile, situational context, and a stream of new information. The model will be evaluated based on the produced

system-initiated conversations (if needed). Having a single evaluation metric that can reflect all aspect of conversation initiation evaluation would be challenging and require further investigation. Approaches like economic models of interactive information retrieval that model the system by assigning cost and benefit to each interaction may be relevant. In case of online evaluation, the typical A/B tests can be used to evaluate the system, and the system can be evaluated by interpreting the positive and negative feedback received from the user. Such feedback can be obtained by identifying the user response type (see Section 4).

6.5. End-to-End Evaluation of Mixed-Initiative Conversations

System-initiated conversations are just one type of interactions that a mixed-initiative CIS system may support. There exist several other interactions, such as the typical user-initiated information seeking conversations and asking clarifying questions for intent disambiguation [21]. The ultimate evaluation methodology should assess the quality of the system in all of these different settings. Such complex end-to-end evaluation can be again done using both online and offline evaluation using a data that contains all different sorts of interactions. Similar approaches as the one mentioned in the last subsection can be adopted, however, designing an evaluation metric for this purpose would be even more challenging.

7. Dangers of System-Initiative Interactions

Privacy Concerns Even with existing conversational information systems, users often have privacy concerns about how their information is processed, to whom it is disclosed and what is the associated risk [22]. We envision that those concerns will only be exacerbated, if left unaddressed, with a system capable of processing far more sensitive user information and engaging in an active form of conversation. Hence we believe that certain privacy concerns must be addressed while designing and implementing active CIS systems. Secure information retrieval and data sharing protocols would be needed to safeguard and ensure users that their identifiable information remains secure. Ensuring and safeguarding user information may or may not instill a sense of security among the end users if the activity format of the underlying system comes off as too intrusive. For instance, one of the use cases for an active engagement CIS is that of *contributing to a multi-party human conversation* (Table 1). Active system engagement in a multi-party setting has been a largely unexplored area in the IR literature and raises new and unique privacy concerns of its

own. For example, one lingering question could be, how should we derive user consent of all the humans involved whose data the system processes? Fully studying privacy implications of such a system would require extensive user studies, often on a task-by-task basis, and assessing perception of the system behavior itself on the end-users.

Badly Timed Engagements Arguably one of the most important components of an active engagement CIS would be its initiator decision making system that decides when to initiate a conversation and perhaps more importantly, when *not* to initiate one. Engagements made by the system at a bad time can be counter-productive or even downright dangerous. For example, while initiating a non time-sensitive conversation, the agent must not disturb or distract the user in any way. An unexpected system engagement when the user is engaged in a critical activity, e.g. driving, can be extremely dangerous and distracting. Therefore, while developing an initiator module we must also account for actively penalizing the module if it engages at a particularly bad time.

8. Broader Impact

This paper highlighted various real-world applications of system-initiative interactions in conversational information seeking systems. The authors believe that research progress in modeling and evaluating system-initiative CIS can potentially lead to a broader impact. Several health and safety related conversations can be initiated by CIS systems to warn users of potential harms and hazards. Such system-initiative interactions can be triggered based on the user's situational context, such as location or health-related signals captured by various sensors embedded into smartphones and wearable devices. Furthermore, these systems can potentially inform the victims of misinformation or abusive content which targets the users through human conversations, written documents, or ads. Different types of entertainment can also be an application of system-initiative interactions, which can be or not be related to information seeking.

Moreover, with the progress of virtual and augmented reality devices, system-initiative interactions (especially those with the information seeking nature) would be of great importance, since the user can experience a virtual environment and a system-initiative CIS can guide the users as they are exploring the virtual environment.

9. Related Work

The study of interaction has a long history in information retrieval research, starting in the 1960s [23]. Much of the earlier research studied how users interacted with intermediaries during information seeking dialogues but this

rapidly shifted to studying how users interacted with operational retrieval systems, including proposals for how to improve the interaction. Information retrieval systems based on this research were also implemented. Oddy [24] developed an interactive information retrieval system with rule-based dialogue interactions in 1977. Croft and Thompson [25] later proposed the first interactive information retrieval system that models user, I³R, using a mixture of expert architecture. A few years later [26] characterized information seeking strategies for interactive IR, offering users choices in a search session based on case-based reasoning.

Since the development of web search engines, research has mostly focused heavily on understanding user interaction with search engines based on an analysis of the search logs available to commercial search engine providers. Since then, explicit modeling of information seeking dialogues or conversations with the aim of improving the effectiveness of retrieval has not been a focus of research until recently. One exception is the TREC Session Track [27] that focused on the development of query formulation during a search session and improving retrieval performance by incorporating knowledge of the session context. On the other hand, commercial personal assistants such as Apple Siri and Google Assistant have become commonplace and there is a clear incentive to develop better conversational models for search. A promising development has been the effectiveness of neural models for generating conversational responses when trained on large amounts data (e.g., [28]).

In recent years, conversational information seeking systems have attracted attention in both academia and the industry [1, 29]. They include conversational search, recommendation, and question answering systems. CIS systems are sufficiently broad to cover a wide range of tasks. The research community has so far studied a number of them, including conversational answer retrieval [2], conversational answer extraction (often referred to as conversational question answering) [3], conversational query re-writing [30], next question prediction [31], speech-only interfaces for conversational search [32], and question-based recommendation (often referred to as conversational recommendation) [33]. In all of these tasks, the user initiates the conversation with the CIS system and the system responds. Even in case of existing conversational recommender systems, the conversations are initiated by users [33]. In this work, however, we discuss challenges and possible solutions for extending existing models to support system-initiative conversations.

System-initiative conversations are indeed related to mixed-initiative interactions (Section 9.1). There are some other related research directions that may be outside of the IR community, including dialogue acts (Section 9.2), system-initiative dialogue systems (Section 9.3)

and push notifications in desktop and mobile apps (Section 9.6). In the following, we present an overview of these related domains to position our work in context.

9.1. Mixed-Initiative Interactions

Most approaches to human-computer interactions with intelligent systems are either controlled by human or system. However, developing intelligent systems that support *mixed-initiative interactions* has always been desired. Allen et al. [7] believed that development of mixed-initiative intelligent systems will ultimately revolutionize the world of computing. Mixed-initiative interactions in dialogue systems have been explored since the 1980s [34, 35, 36]. Early attempts to build systems that support mixed-initiative interactions include the Look-Out system [37] for scheduling and meeting management in Microsoft Outlook, Clippit⁴ for assisting users in Microsoft Office, and TRIPS [38] for assisting users in problem solving and planning.

Horvitz [37] identified 12 principles that systems with mixed-initiative user interfaces must follow. In summary, mixed-initiative interactions should be taken at the right time in the light of cost, benefit, and uncertainties. Many factors can impact cost and benefit of interactions that are covered in multiple principles. In addition, systems with mixed-initiative interactions should put user at the center and allow efficient invocation and termination. They are expected to memorize past interactions and continuously learn by observation. Based on these principles, conversational systems by nature raise the opportunity of mixed-initiative interactions.

Allen et al. [7] defined four levels of mixed-initiative interactions in the context of dialogue systems, as follows:

1. **Unsolicited reporting:** An agent notifies others of critical information as it arises. For example, an agent may constantly monitor the progress for the plan under development. In this case, the agent can notify the other agents (e.g., user) if the plan changes.
2. **Subdialogue initiation:** An agent initiates subdialogues to clarify, correct, and so on. For example, in a dialogue between a user and a system, the system may ask a question to clarify the user's intent. Since the system asks the question and the user should answer the question, and clarification may take multiple interactions, the system has temporarily taken the initiative until the issue is resolved. This is why it is called subdialogue initiation.
3. **Fixed subtask initiation:** An agent takes initiative to solve predefined subtasks. For example, if an agent is supposed to complete a task that involves multiple

⁴https://en.wikipedia.org/wiki/Office_Assistant

subtasks. In this case, the agent can take initiative to ask questions and complete the subtask. Once the subtask is completed, initiative reverts to the user.

4. **Negotiated mixed-initiative:** Agents coordinate and negotiate with other agents to determine initiative. This is mainly defined for multi-agent systems in which agents decide whether they are qualified to complete a task or it should be left for other agents.

When it comes to open-domain conversational information seeking, some of these mixed-initiative levels remain valid. Mixed-initiative conversational information seeking has relatively less explored, nevertheless identified as critical components of a conversational system [6, 39]. Perhaps clarification [8, 40, 10] and preference elicitation [11, 12] are the two areas related to mixed-initiative interactions that have attracted much attention in recent years. However, they are mostly unrelated to system-initiative interactions, which are relatively unexplored. Nevertheless, the unsolicited reporting level of mixed-initiative interactions mentioned above include several interesting example use-cases for system-initiative CIS systems.

9.2. Dialogue Acts in Conversational Systems

Spoken dialogue systems (SDS) have allowed for interaction with computer-based applications (e.g., smart speakers) through spoken natural language. Certain SDS mechanisms are specifically designed to carry out well-defined tasks, e.g., scheduling [41], and most of them are based on a finite state-based dialogue control. Although the focus of CIS research is mostly on open-domain information seeking tasks, such dialogue acts can be potentially used to support a diverse set of modes and scenarios in system-initiative CIS systems. A range of prior studies in dialogue acts also offer insights into designing models for conveying information through conversations e.g. prior work by Bunt et al. [42] offers promising features derived from broad dialogues to better model information needs, however in our work we assume that an alternate method might be required to extract similar information-needs from user meta-data. Dialogue acts can potentially serve as communicative functions of dialogue segments, such as *request*, *inform*, *question*, *suggest* and *offer*. The general taxonomy of dialogue acts is complex with different markup schemes. One segment of particular interest to us, and often not examined in the IR literature, is that of *turn taking* [43]. For instance, our conversational agent will have control over the dialogue and the segments might be produced through an analysis of user data. Such an analysis over user meta-data (e.g. Location) to personalize an IR task isn't new, and is most commonly applied in the context of personalized web-searches [44].

Some of the personalization methods leverage long-term user behavioral histories [45], while others analyze short-term *implicit* feedback [46]. A key challenge that we see with personalization, especially when applied to active-engagement conversational systems, is that of collecting user profiles with sufficiently rich features, while balancing privacy concerns. We leave that as an essential component of our future work.

9.3. Initiative Control in Dialogue Systems

Discourse segmentation through transfer of control in dialogue systems was first studied decades ago by Walker and Whittaker [47] to enable flexible human-computer conversations to take place that allow for corrections and clarifications. Since then a number of studies have been done to determine the *ideal* behavior of a virtual assistant [48, 49]. One of the key aspects of such an ideal behavior is *initiative*. In prior work, a number of authors have considered what constitutes initiative [50, 51, 52]. Instead of a human, at certain relevant points in time, the system may take the initiative to engage in a conversation. Among all such systems, there is some form of *dialogue management* component which determines what to prompt for and/or what to accept next based on the conversation history and its context [53]. Such a management component plays a central role in the traditional architecture of a dialogue system and is primarily concerned with the flow of the dialogue (information providing, feedback request, etc) while simultaneously maintaining a discourse history. For instance, Vakulenko et al. [54] have shown how an agent might effectively take initiative to elicit or clarify information when appropriate.

In addition to standalone engagement by a conversational agent, studies have also shown that sources of information that led to that engagement are equally important – e.g., different sources have varying influence on purchase decisions, implying that the effectiveness of a conversational information system depends on the system saying *why* it made a specific decision or recommendation [28, 55].

9.4. Contextual Suggestions

Contextual suggestion track within TREC [56] is aimed at providing personalized point-of-interest recommendations to users in a ranked manner. The task assumes a certain setting – a user in a specific place (geographic location) with a trip type. Given the same user's personal profile (interests, endorsements etc), the system makes recommendations for attractions. The track consisted of two phases, *Phase 1*: participants could select any venue from the reference collection. *Phase 2*: participants had to

rank a given list of venues for each user and thus allowing for ground truth data against which the system could be evaluated. Early works on this task involved rule-based approaches by mapping user-profiles to specific venues. Recently, people have experimented with standard machine learning [57] and neural methods [58, 59] for best mapping user profiles to relevance-rated documents. For example, Seyler et al. [58] create graph embeddings from a heterogeneous information network (HIN) using the TREC Contextual Suggestion dataset achieving state-of-the-art performance.

9.5. Incident Streams

TREC-IS track in 2018 [60] focused on curating feeds of social media posts and classify them based on *actionable information* for enhanced situational awareness (such as emergencies). Incident streams are relevant in context of system-initiative conversations due to the underlying nature of the task – analyze large sets of textual information related to user profiles and act in a time-sensitive manner. In addition to the type of information, TREC-IS evaluation tasks also include a criticality-score indicating how *important* it is for a specific content to be acted upon.

9.6. Push Notifications

Push notifications have been mostly studied in the context of mobile applications, largely with e-commerce goals [61, 62, 63]. Much of the prior research on push notifications has focused on their disruptive nature [64, 65]. For example, Mehrotra et al. [66] provided an in-depth study evaluating how the user-response time of a non time sensitive notification is influenced by the notification’s presentation, modality as well as the sender-recipient relationship. Mehrotra et al. [67] further detailed, through extensive user studies, how push notifications with different context and timings can cause disruptions. This is an especially important component since one of the main goals of an active engagement system is to minimize disruptions caused to the end-user. Other works in the area have explored the use of push notifications for meta-learning [68] and self-logging [69] to better adapt the underlying framework for adjusting user preferences. Push notifications are basically system-initiative interactions, however they mostly do not concern with information seeking tasks and are fundamentally different from system-initiative interactions in conversational systems.

9.7. Information Need in Collaborative Conversations

Over time, a number of definitions for *information need* have been conceptualized [70, 71]. For our work, we consider the one by Case [72] i.e. *information need is a recognition that the user’s knowledge is inadequate to satisfy their own goals*, as it implies that the *information need* must emerge from the user’s end. Collaborative conversations offer one such instance, as articulated by Shiga et al. [73], that information needs in such conversations are naturally verbalized and therefore can be captured by end-user devices. Furthermore, we can utilize the taxonomy of information needs defined by Taylor [74] to differentiate between perceived needs and actual queries since Taylor’s model consists of *visceral, conscious, formalized and compromised needs*. For the purposes of a conversational information system to actively engage in a collaborative discussion, we primarily focus on the *conscious needs*, which are defined as “ambiguous and rambling statement” but ultimately evolve into *formalized needs* (qualified and rational). Prior work by Jansen et al. [75] on analyzing conversation query logs has shown that users often frame a short and under-specified query to information seeking systems, however community-based modern QA models are often capable to *formalizing* such information needs on QA sites or speech-oriented search systems. The degree of interest in collaborative conversational information search has increased since then and has led to quantitative analysis of conversations during search [76, 77, 78]. For example, Foster [79] performed a full discourse analysis on group conversations to define the relationship between *functions* of verbal context and information seeking activity. While these studies remain either conceptual or use small amounts of textual chat data, they nevertheless suggest that collaborative conversations can be a useful source for conversational information seeking.

In this research, we also highlight some applications of system-initiative conversational interactions in the context of multi-party and collaborative conversations.

10. Summary

In this work, we explored applications and the ways to model an active engagement conversational information seeking system. We defined a taxonomy upon which a framework for an active engagement system could be built. Our taxonomy defines three broad dimensions of an active engagement framework – initiation moment (*when to initiate a conversation*), initiation purpose (*why to initiate a conversation*) and interaction means (*how to initiate a conversation*). Subsequently we show, through the explained examples and a pipeline, how the described

characteristics are both necessary and sufficient to allow for the functioning of an active engagement information seeking system, for a number of initiation purposes. In doing so, we also generalized several components of the pipeline that have been implemented before with proven effectiveness. We view the contribution of our work as this taxonomy and the proposed generic pipeline, which can be employed towards building true active-engagement information seeking systems. Implementing and evaluating the proposed framework in a user-centric way remains the most important future directions suggested by this work. It is further worth considering the numerous technical and evaluation challenges that come with the proposed approach. Finally, we believe that our identified use cases of active engagement CIS systems would only serve as founding basis for other, broader case-specific applications such as aiding people with disabilities, as highlighted in section 8 and we hope this work spurs additional work in this largely-unexplored area of information seeking research.

Acknowledgments

This work was supported in part by the Center for Intelligent Information Retrieval. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

References

- [1] J. S. Culpepper, F. Diaz, M. D. Smucker, Research frontiers in information retrieval: Report from the third strategic workshop on information retrieval in lorne (swirl 2018), SIGIR Forum 52 (2018) 34–90.
- [2] C. Qu, L. Yang, C. Chen, M. Qiu, W. B. Croft, M. Iyyer, Open-Retrieval Conversational Question Answering, SIGIR '20, ACM, New York, NY, USA, 2020, p. 539–548.
- [3] S. Reddy, D. Chen, C. D. Manning, CoQA: A conversational question answering challenge, Transactions of the Association for Computational Linguistics 7 (2019) 249–266.
- [4] Y. Sun, Y. Zhang, Conversational recommender system, in: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18, ACM, New York, NY, USA, 2018, p. 235–244.
- [5] P. Thomas, D. McDuff, M. Czerwinski, N. Craswell, Misc: A data set of information-seeking conversations, in: Proceedings of the 1st International Workshop on Conversational Approaches to Information Retrieval, CAIR '17, 2017.
- [6] F. Radlinski, N. Craswell, A theoretical framework for conversational search, in: Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, CHIIR '17, ACM, New York, NY, USA, 2017, p. 117–126.
- [7] J. E. Allen, C. I. Guinn, E. Horvitz, Mixed-initiative interaction, IEEE Intelligent Systems and their Applications 14 (1999) 14–23.
- [8] M. Aliannejadi, H. Zamani, F. Crestani, W. B. Croft, 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 '19, ACM, New York, NY, USA, 2019, p. 475–484.
- [9] H. Zamani, M. Bendersky, X. Wang, M. Zhang, Situational context for ranking in personal search, in: Proceedings of the 26th International Conference on World Wide Web, WWW '17, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2017, p. 1531–1540.
- [10] H. Zamani, B. Mitra, E. Chen, G. Lueck, F. Diaz, P. N. Bennett, N. Craswell, S. T. Dumais, Analyzing and learning from user interactions for search clarification, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, ACM, New York, NY, USA, 2020, p. 1181–1190.
- [11] F. Radlinski, K. Balog, B. Byrne, K. Krishnamoorthi, Coached conversational preference elicitation: A case study in understanding movie preferences, in: Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue, ACL, Stockholm, Sweden, 2019, pp. 353–360.
- [12] A. Seplarskaia, J. Kiseleva, F. Radlinski, M. de Rijke, Preference elicitation as an optimization problem, in: Proceedings of the 12th ACM Conference on Recommender Systems, RecSys '18, ACM, New York, NY, USA, 2018, p. 172–180.
- [13] N. J. Belkin, W. B. Croft, Information filtering and information retrieval: Two sides of the same coin?, Commun. ACM 35 (1992) 29–38.
- [14] S. Robertson, D. A. Hull, The TREC-9 Filtering Track Final Report, in: Proceedings of Ninth Text Retrieval Conference, TREC-9, National Institute of Standards and Technology, Special Publication, 2000, pp. 25–40.
- [15] A. Weiss, R. Bernhaupt, D. Dürnberger, M. Altmaninger, R. Buchner, M. Tscheligi, User experience evaluation with a wizard of oz approach: Technical and methodological considerations, 2010, pp. 303 – 308.
- [16] K. Järvelin, J. Kekäläinen, Cumulated gain-based evaluation of ir techniques, volume 20, ACM, New

- York, NY, USA, 2002, p. 422–446.
- [17] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, Bleu: a method for automatic evaluation of machine translation, in: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, ACL, Philadelphia, Pennsylvania, USA, 2002, pp. 311–318.
- [18] C.-Y. Lin, ROUGE: A package for automatic evaluation of summaries, in: Text Summarization Branches Out, ACL, Barcelona, Spain, 2004, pp. 74–81.
- [19] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, Y. Artzi, Bertscore: Evaluating text generation with bert, ArXiv abs/1904.09675 (2020).
- [20] N. Mathur, T. Baldwin, T. Cohn, Tangled up in BLEU: Reevaluating the evaluation of automatic machine translation evaluation metrics, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL, Online, 2020, pp. 4984–4997.
- [21] E. Pitler, K. Church, Using word-sense disambiguation methods to classify web queries by intent, in: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP '09, ACL, Singapore, 2009, pp. 1428–1436.
- [22] S. Zimmerman, A. Thorpe, C. Fox, U. Kruschwitz, Investigating the interplay between searchers' privacy concerns and their search behavior, in: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '19, ACM, New York, NY, USA, 2019, p. 953–956.
- [23] D. Kelly, C. R. Sugimoto, A systematic review of interactive information retrieval evaluation studies, 1967-2006, *J. Assoc. Inf. Sci. Technol.* 64 (2013) 745–770.
- [24] R. Oddy, Information retrieval through man-machine dialogue, *Journal of Documentation* 33 (1977) 1–14.
- [25] W. B. Croft, R. H. Thompson, I³r: A new approach to the design of document retrieval systems, *JASIS* 38 (1987) 389–404.
- [26] N. J. Belkin, C. Cool, A. Stein, U. Thiel, Cases, scripts, and information-seeking strategies: On the design of interactive information retrieval systems, *Expert Systems with Applications* 9 (1995) 379–395.
- [27] B. Carterette, P. Clough, M. Hall, E. Kanoulas, M. Sanderson, Evaluating retrieval over sessions: The trec session track 2011-2014, in: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16, ACM, New York, NY, USA, 2016, pp. 685–688.
- [28] A. Sordani, M. Galley, M. Auli, C. Brockett, Y. Ji, M. Mitchell, J.-Y. Nie, J. Gao, B. Dolan, A neural network approach to context-sensitive generation of conversational responses, in: Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL '15, ACL, Denver, Colorado, 2015, pp. 196–205.
- [29] A. Anand, L. Cavedon, H. Joho, M. Sanderson, B. Stein, Conversational search (dagstuhl seminar 19461), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [30] S. Yu, J. Liu, J. Yang, C. Xiong, P. Bennett, J. Gao, Z. Liu, Few-shot generative conversational query rewriting, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, ACM, New York, NY, USA, 2020, p. 1933–1936.
- [31] L. Yang, H. Zamani, Y. Zhang, J. Guo, W. Croft, Neural matching models for question retrieval and next question prediction in conversation, in: Proceedings of the 2017 ACM SIGIR Workshop on Neural Information Retrieval, NeuIR @ SIGIR '17, 2017.
- [32] J. R. Trippas, Spoken conversational search: Speech-only interactive information retrieval, in: Proceedings of the 2016 ACM Conference on Human Information Interaction and Retrieval, CHIIR '16, ACM, New York, NY, USA, 2016, p. 373–375.
- [33] W. Lei, X. He, M. de Rijke, T.-S. Chua, Conversational Recommendation: Formulation, Methods, and Evaluation, SIGIR '20, ACM, New York, NY, USA, 2020, p. 2425–2428.
- [34] H. Kitano, C. Van Ess-Dykema, Toward a plan-based understanding model for mixed-initiative dialogues, in: Proceedings of the 29th Annual Meeting on Association for Computational Linguistics, ACL '91, ACL, USA, 1991, p. 25–32.
- [35] D. G. Novick, S. A. Douglas, Control of Mixed-Initiative Discourse through Meta-Locutionary Acts: A Computational Model, Technical Report, USA, 1988. AAI8911322.
- [36] M. Walker, S. Whittaker, Mixed initiative in dialogue: An investigation into discourse segmentation, in: Proceedings of the 28th Annual Meeting on Association for Computational Linguistics, ACL '90, ACL, USA, 1990, p. 70–78.
- [37] E. Horvitz, Principles of mixed-initiative user interfaces, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '99, ACM, New York, NY, USA, 1999, pp. 159–166.
- [38] G. Ferguson, J. F. Allen, TRIPS: an integrated intelligent problem-solving assistant, in: Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI '98, IAAI '98, AAAI Press / The MIT Press, 1998, pp. 567–572.
- [39] J. R. Trippas, D. Spina, L. Cavedon, H. Joho,

- M. Sanderson, Informing the design of spoken conversational search: Perspective paper, in: Proceedings of the 2018 Conference on Human Information Interaction & Retrieval, CHIIR '18, ACM, New York, NY, USA, 2018, p. 32–41.
- [40] H. Zamani, S. Dumais, N. Craswell, P. Bennett, G. Lueck, Generating clarifying questions for information retrieval, in: Proceedings of The Web Conference 2020, WWW '20, ACM, New York, NY, USA, 2020, p. 418–428.
- [41] M. McTear, Z. Callejas, D. Griol, *The Conversational Interface: Talking to Smart Devices*, 1st ed., Springer Publishing Company, Incorporated, 2016.
- [42] H. Bunt, J. Alexandersson, J. Carletta, J.-W. Choe, A. C. Fang, K. Hasida, K. Lee, V. Petukhova, A. Popescu-Belis, L. Romary, C. Soria, D. Traum, Towards an ISO standard for dialogue act annotation, in: Proceedings of the Seventh International Conference on Language Resources and Evaluation, LREC '10, ELRA, Valletta, Malta, 2010.
- [43] H. Khouzaimi, R. Laroche, F. Lefèvre, Turn-taking phenomena in incremental dialogue systems, in: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP '15, ACL, Lisbon, Portugal, 2015, pp. 1890–1895.
- [44] P. N. Bennett, F. Radlinski, R. W. White, E. Yilmaz, Inferring and using location metadata to personalize web search, in: Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11, ACM, New York, NY, USA, 2011, p. 135–144.
- [45] N. Matthijs, F. Radlinski, Personalizing web search using long term browsing history, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11, ACM, New York, NY, USA, 2011, p. 25–34.
- [46] X. Shen, B. Tan, C. Zhai, Context-sensitive information retrieval using implicit feedback, in: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05, ACM, New York, NY, USA, 2005, p. 43–50.
- [47] M. Walker, S. Whittaker, Mixed initiative in dialogue: An investigation into discourse segmentation, in: Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics, ACL '90, ACL, Pittsburgh, Pennsylvania, USA, 1990, pp. 70–78.
- [48] J. Cassell, Embodied conversational interface agents, *Commun. ACM* 43 (2000) 70–78.
- [49] Z. Wang, J. Lee, S. Marsella, Towards more comprehensive listening behavior: Beyond the bobble head, *Intelligent Virtual Agents* (2011) 216–227.
- [50] J. E. Allen, C. I. Guinn, E. Horvitz, Mixed-initiative interaction, *IEEE Intelligent Systems and their Applications* 14 (1999) 14–23.
- [51] J. Chu-Carroll, M. K. Brown, An evidential model for tracking initiative in collaborative dialogue interactions, *User Model. User Adapt. Interact.* 8 (1998) 215–254.
- [52] Z. Wang, J. Lee, S. Marsella, Towards more comprehensive listening behavior: Beyond the bobble head, in: *Intelligent Virtual Agents*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 216–227.
- [53] C. Lee, S. Jung, K. Kim, D. Lee, G. Lee, Recent approaches to dialog management for spoken dialog systems, *J. Comput. Sci. Eng.* 4 (2010) 1–22.
- [54] S. Vakulenko, E. Kanoulas, M. de Rijke, An Analysis of Mixed Initiative and Collaboration in Information-Seeking Dialogues, *SIGIR '20*, ACM, New York, NY, USA, 2020, p. 2085–2088.
- [55] N. Tintarev, J. Masthoff, Effective explanations of recommendations: User-centered design, in: Proceedings of the 2007 ACM Conference on Recommender Systems, RecSys '07, ACM, New York, NY, USA, 2007, p. 153–156.
- [56] S. H. Hashemi, J. Kamps, J. Kiseleva, C. L. Clarke, E. M. Voorhees, Overview of the trec 2016 contextual suggestion track., in: *TREC*, 2016.
- [57] M. Aliannejadi, F. Crestani, Venue appropriateness prediction for personalized context-aware venue suggestion, in: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 1177–1180.
- [58] D. Seyler, P. Chandar, M. Davis, An information retrieval framework for contextual suggestion based on heterogeneous information network embeddings, in: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 953–956.
- [59] J. Mo, L. Lamontagne, R. Khoury, Word embeddings and global preference for contextual suggestion, in: *TREC*, 2016.
- [60] R. McCreddie, C. Buntain, I. Soboroff, Trec incident streams: Finding actionable information on social media, in: *ISCRAM*, 2019.
- [61] A. Kumar, S. Johari, Push notification as a business enhancement technique for e-commerce, in: Proceedings of the Third International Conference on Image Information Processing, ICIIP '15, 2015, pp. 450–454.
- [62] V. S. Moertini, C. D. Nugroho, E-commerce mobile marketing model resolving users acceptance criteria, *International Journal of Managing Information Technology* 4 (2012) 23–40.
- [63] L. Tan, A. Roegiest, J. Lin, C. L. Clarke, An explo-

- ration of evaluation metrics for mobile push notifications, in: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16, ACM, New York, NY, USA, 2016, p. 741–744.
- [64] A. Sahami Shirazi, N. Henze, T. Dingler, M. Pielot, D. Weber, A. Schmidt, Large-scale assessment of mobile notifications, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '14, ACM, New York, NY, USA, 2014, p. 3055–3064.
- [65] M. Pielot, K. Church, R. de Oliveira, An in-situ study of mobile phone notifications, in: Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services, MobileHCI '14, ACM, New York, NY, USA, 2014, p. 233–242.
- [66] A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley, M. Musolesi, My phone and me: Understanding people's receptivity to mobile notifications, in: Proceedings of the 2016 SIGCHI Conference on Human Factors in Computing Systems, CHI '16, ACM, New York, NY, USA, 2016, p. 1021–1032.
- [67] A. Mehrotra, R. Hendley, M. Musolesi, Prefminer: Mining user's preferences for intelligent mobile notification management, in: Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '16, ACM, New York, NY, USA, 2016, p. 1223–1234.
- [68] B. Tabuenca, M. Kalz, S. Ternier, M. Specht, Stop and think: Exploring mobile notifications to foster reflective practice on meta-learning, IEEE Transactions on Learning Technologies 8 (2015) 124–135.
- [69] F. Bentley, K. Tollmar, The Power of Mobile Notifications to Increase Wellbeing Logging Behavior, CHI '13, ACM, New York, NY, USA, 2013, p. 1095–1098.
- [70] S. M. Gray, Looking for information: A survey of research on information seeking, needs, and behavior, Journal of the Medical Library Association 91 (2003) 259–260.
- [71] A. Stolcke, K. Ries, N. Coccaro, E. Shriberg, R. Bates, D. Jurafsky, P. Taylor, R. Martin, C. Van Ess-Dykema, M. Meteer, Dialogue act modeling for automatic tagging and recognition of conversational speech, Computational Linguistics 26 (2000) 339–374.
- [72] D. Case, Looking for information—a survey of research on information seeking, needs, and behavior, Inf. Res. 8 (2003).
- [73] S. Shiga, H. Joho, R. Blanco, J. R. Trippas, M. Sanderson, Modelling information needs in collaborative search conversations, in: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17, ACM, New York, NY, USA, 2017, p. 715–724.
- [74] R. S. Taylor, The process of asking questions, American Documentation 13 (1962) 391–396.
- [75] B. J. Jansen, A. Spink, T. Saracevic, Real life, real users, and real needs: a study and analysis of user queries on the web, Information Processing & Management 36 (2000) 207 – 227.
- [76] R. Kelly, S. J. Payne, Collaborative web search in context: A study of tool use in everyday tasks, in: Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '14, ACM, New York, NY, USA, 2014, p. 807–819.
- [77] C. I. Seeking, Collaborative information seeking - best practices, new domains and new thoughts | preben hansen | springer, 2015.
- [78] C. Shah, R. González-Ibáñez, Exploring information seeking processes in collaborative search tasks, in: Proceedings of the 73rd ASIS&T Annual Meeting on Navigating Streams in an Information Ecosystem - Volume 47, ASIS&T '10, American Society for Information Science, USA, 2010.
- [79] J. Foster, Understanding interaction in information seeking and use as a discourse: a dialogic approach, J. Documentation 65 (2009) 83–105.