Epistemo: A Crowd-Powered Conversational Search Interface

Saiganesh Swaminathan
CMU
Pittsburgh, PA 15217, USA
saiganes@cs.cmu.edu

Ting-Hao (Kenneth) Huang
CMU
Pittsburgh, PA 15217, USA
tinghaoh@andrew.cmu.edu

Irene Lin
CMU
Pittsburgh, PA 15217, USA
iwl@andrew.cmu.edu

Anhong Guo
CMU
Pittsburgh, PA 15217, USA
anhongg@cs.cmu.edu

Gierad Laput
CMU
Pittsburgh, PA 15217, USA
gierad.laput@cs.cmu.edu

Jeffrey Bigham
CMU
Pittsburgh, PA 15217, USA
jbigham@cs.cmu.edu

Abstract
Conversational assistants have the potential to help people find and make sense of personal information as they go about their everyday lives. We are focusing on email because it is ubiquitous and often serves the function of an extension of our memory, containing our reservation numbers, address, images, phone numbers, personal documents, etc. Yet, support for conversational interaction is limited, both by the brittleness of current automated approaches and by privacy concerns introduced when considering private repositories like email.

Epistemo is our conversational search system that enables users to search, extract and manage information from emails in a privacy-preserving way by having crowds write programs on-the-fly to help the automated system work better. We first describe the design of our system, how it works and later discuss social implications arising from its design. Our contributions to the workshop are thus 1) to initiate a discussion on how crowd powered conversational search interfaces can be better designed and 2) to gather inputs on what may be the gaps that needs to filled to take such systems to real world.

Author keywords
Conversation, search, email, collaboration, personal information management, crowd-powered system

ACM Classification Keywords
H.5.0. Information interfaces and presentation (e.g., HCI): General

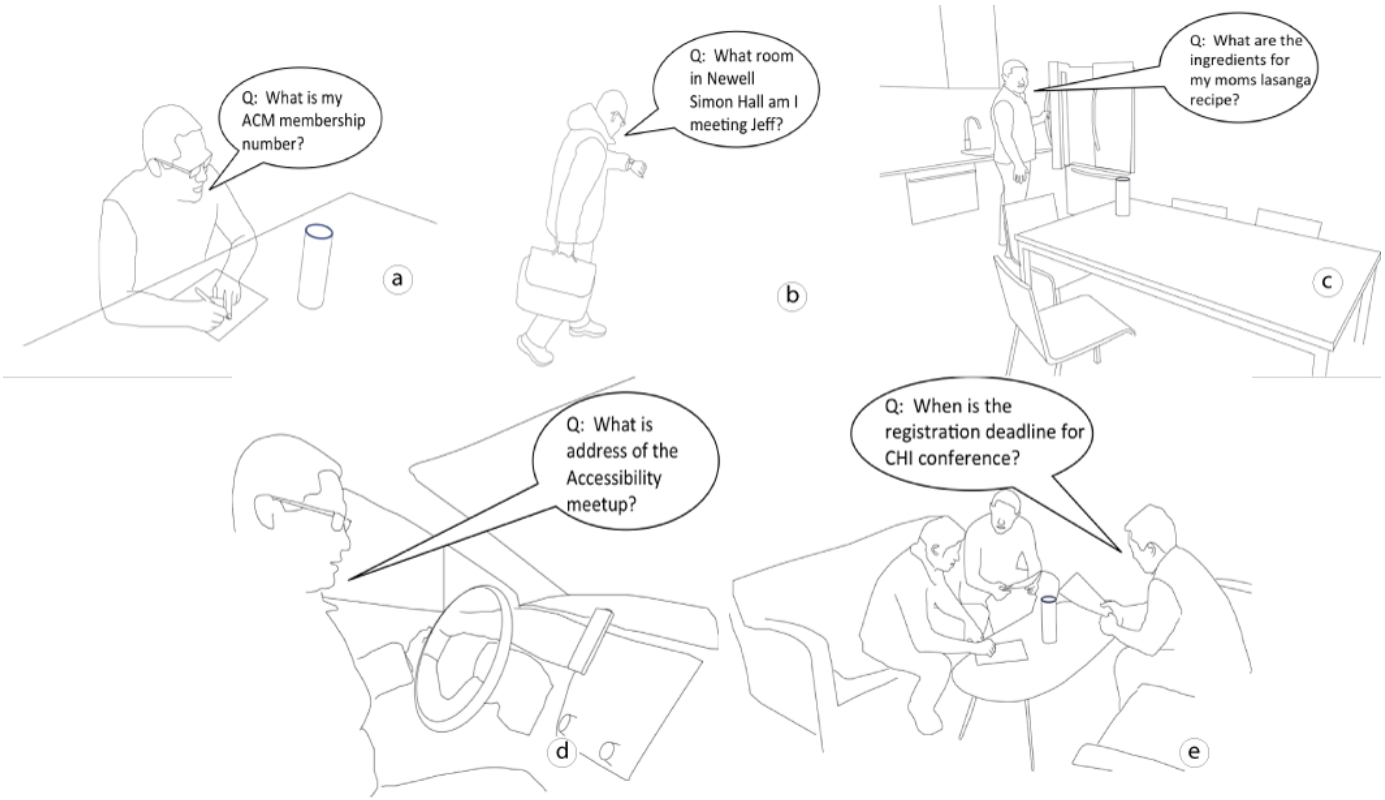

Figure 1: Examples of tasks we envision using Epistemo for: (a) Asking for ACM membership when filling a form (b) Asking for room number on the way to a meeting (c) Asking for ingredients of lasagna recipe your mom sent you (d) Asking for address of a accessibility meetup (e) Collaboratively searching for which conferences might be suitable for the joint paper. Each conversational task is made more robust via the on-demand assistance of an online crowd, and privacy is maintained by never showing the email text directly to crowd workers but rather by having them write programs to help the automated system work better.

Introduction and Related work

Conversational Search offers new ways to help users satisfy their everyday information needs. However their capabilities are limited to what the system developers have designed them for and doesn't often meet users expectations [1].

To expand search capabilities beyond the popular information needs [2], we propose a novel crowd-powered approach. We explore our approach through Epistemo – a conversational search interface which helps users search and manage information contained in emails. As such, emails serves as our extended memory containing information pertaining to everyday life needs (see fig1) and searching for just the information we want remains laborious. The problem of retrieving information from emails has been studied through systems like *Phlat* [4] and *Stuff I have seen* [3] where users search for personal information through interactive queries of metadata.

Such tools offer help to find information in personal digital archives, however the search process is quite demanding as it consists of many steps (formulating and re-formulating queries, scanning through search results, etc.) to be performed individually by the user. Hence, our approach instead is to involve crowd in conversational manner to help users in every step of the search process *i.e,* filtering, ranking and retrieving information from emails.

Related work
We build on previous related work on using crowd to manage emails [5], helping people extract information from the web using crowd[6], performing query expansion with the crowd [7] and finally on building privacy sensitive crowdsourcing workflows [8].

System and crowd workflow

We now go over our crowd work flow and explain how crowd can help in each step of the conversational search process.

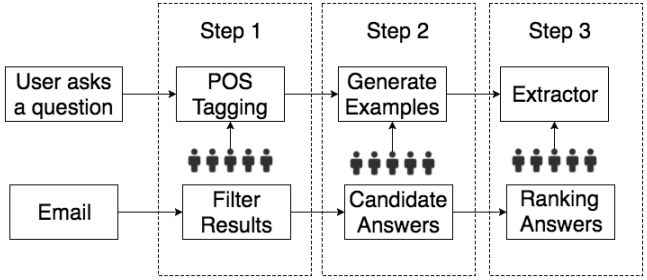

Step #1: *Crowd constructs query from conversation*
In this stage the users questions first gets tagged into a query for searching through their emails. We extract the noun Phrase from each question using a syntactic parser and then use it as a query to search emails. To increase our precision, we ask the crowd if more words need to be included from the question into the query.

Step #2: *Epistemo collects examples or input for Named Entity Recognition*
After retrieving the emails using query generated from step1, we now have to extract the information the user is looking for in the emails. To achieve this, we ask the crowd to select the type of entity being asked by the user (*e.g*, Time, Date, Money, Location, etc). We then use workers input to feed a named entity recognizer[10] which loops through the retrieved emails to extract the information.

Alternatively, If an entity being asked by the user does not exist, we ask crowd workers to fetch examples of how they look like from the internet *e.g,* ACM membership number, building numbers, etc. Based on these examples, we create a simple extractor using regular expression and run it over our retrieved emails.

Step #3: *Crowd moderating search results and user queries again*
Through step1 and step2, we are able to extract a list of possible candidate answers. If the user is not satisfied with the retrieved information, we involve the crowd to re-rank the search results by obfuscating privacy sensitive parts of the retrieved email. This process continues until user gets the answer he seeks.

Future Work

Several interesting social implications remain to be explored by the design of such crowd powered conversational search systems. We hope to motivate the audience at the workshop about implications of designing such a system and will elicit ideas on how we can improve. Here is a list of few exemplar:

1. Asking questions and Navigating search results
Current conversational search interfaces have limited affordances in understanding what the users want, to that extent that we have books teaching us, how to talk to them [9]. Through Epistemo, we aim to bridge this gap by involving humans in the loop, However this introduces several new social implications

a. How will users talk or formulate questions to a conversational search interfaces knowing humans are involved in the loop? Huang et al [11] have started investigating the effect of such system very recently, *e.g.,* Will they formulate queries to include more polite words like please, thank you, etc.
b. How long are they willing to spend in the search process through conversations? The average conversation length is currently at 10min [11]. Thus,, the process of finding the right information may take a few iterations, *i.e,* speed vs quality

2. Privacy and Ethical challenges
In the current design of our system, we provide a privacy preserving approach *i.e,* the crowd workers never have direct access to users emails. They simply find examples from the internet on how patterns of information look like and our system takes as input these patterns. There are interesting questions on how much the users are able to trust the system? Does a crowd-powered conversational search interface change the type of questions being asked by the user, *e.g.*, "What is my ACM Membership number" vs "What is my credit card number?"

2. Multi-user conversational search
If there are multiple users involved in searching, how can we design systems to enable turn taking and authorize users to search in each others emails. *e.g,* In fig1(e), we see multiple authors collaborating on a project, they are collaboratively trying to figure out suitable conference, each user has a list of conference and dates in their emails. How can they share their content to enable conversational search

References

1. Luger, Ewa, and Abigail Sellen. "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, 2016.
2. Email search from google, url: https://support.google.com/websearch/answer/1710607?ref_topic=3036132&hl=en

3. Dumais, Susan, Edward Cutrell, Jonathan J. Cadiz, Gavin Jancke, Raman Sarin, and Daniel C. Robbins. "Stuff I've seen: a system for personal information retrieval and re-use." In ACM SIGIR Forum, vol. 49, no. 2, pp. 28-35. ACM, 2016. Harvard

4. Cutrell, Edward, Daniel Robbins, Susan Dumais, and Raman Sarin. "Fast, flexible filtering with phlat." In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pp. 261-270. ACM, 2006.

5. Kokkalis, Nicolas, Thomas Köhn, Carl Pfeiffer, Dima Chornyi, Michael S. Bernstein, and Scott R. Klemmer. "EmailValet: Managing email overload through private, accountable crowdsourcing." In *Proceedings of the 2013 conference on Computer supported cooperative work*, pp. 1291-1300. ACM, 2013.

6. Bernstein, Michael S., Jaime Teevan, Susan Dumais, Daniel Liebling, and Eric Horvitz. "Direct answers for search queries in the long tail." In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 237-246. ACM, 2012.

7. Parameswaran, Aditya, Ming Han Teh, Hector Garcia-Molina, and Jennifer Widom. "Datasift: An expressive and accurate crowd-powered search toolkit." In *First AAAI Conference on Human Computation and Crowdsourcing*. 2013.

8. Lasecki, Walter S., Jaime Teevan, and Ece Kamar. "Information extraction and manipulation threats in crowd-powered systems." In Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing, pp. 248-256. ACM, 2014.

9. Sadun, Erica, and Steve Sande. Talking to Siri: Mastering the Language of Apple's Intelligent Assistant. Que Publishing, 2014. Harvard

10. Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. Proceedings of the 43nd Annual Meeting of the Association for Computational Linguistics (ACL 2005), pp. 363-370.

11. Huang, Ting-Hao Kenneth, Walter S. Lasecki, Amos Azaria, and Jeffrey P. Bigham. ""Is there anything else I can help you with?": Challenges in Deploying an On-Demand Crowd-Powered Conversational Agent." (2016). Harvard