A Context/Communication Information Agent

Jason I. Hong and James A. Landay

Group for User Interface Research, Computer Science Division University of California, Berkeley, CA, USA

Abstract: We are developing a Context/Communication Information Agent (CIA), an autonomous software agent that proactively searches for the right information at the right time. Our goal is to design and evaluate a system that leverages what people naturally do, using this knowledge to retrieve information, and presenting it with a minimal cost of disruption to the users. In this paper, we describe the results a low-fidelity prototype performed in a meeting situation, the design space for such an application, and our plans for continued investigation.

Keywords: Attention; Communication; Context; Information retrieval agent

1. Introduction

The system we envision is a proactive software agent that uses context and human-to-human communication to help find and deliver the right information at the right time. The system constantly searches for information related to the current situation, in order to make it easier to find relevant related information. We call such a system a Context/Communication Information Agent (CIA). In most cases, the agent will keep the information in the background, ready for quick access if the user wants, but in special cases, when the information is especially timely and useful, the agent will call attention to itself.

2. Context and Communication

By context, we mean knowing the answers to the "W" questions, such as *who* is speaking, *who* else is here, *where* am I, *what* calendar event is current, and so on. As an example of how context can be used, suppose that earlier in the day, Francis scribbled down a grocery list on a piece of paper with an electronic pen that records what is written. The handwriting recognition is not perfect, but the spatial structure of the document suggests a list, and enough words are recognised as food items. Later, when passing by his local grocery store, Francis' personal digital assistant (PDA) beeps, reminding him to purchase some groceries. As he enters the store, his CIA fetches his handwritten notes for him and stores them on his PDA. Although Francis remembers what he wrote down earlier for his grocery list, his CIA has made the list easy to access if he wants.

As another example of how context can be used to retrieve information, suppose that Victoria has a weekly staff meeting to attend, stored as a recurring weekly event on her electronic calendar. Every week, at the beginning of the meeting, Victoria always looks for the minutes and action items from last week's meeting. As she searches, the CIA's prediction engine is running in the background, guessing which documents she is looking for. After a few weeks, the prediction engine has, with high confidence, matched Victoria's recurring calendar event to retrieving last week's minutes and action items. When the next meeting takes place, the system begins retrieving the minutes and action items from the previous week's meeting for her, so that she doesn't have to search for where she saved them.

By human-to-human communication, we mean using microphones, cameras and other sensors to capture natural modes of communication between people, such as speaking, writing and sketching. As an example of how communication could be used to fetch information, suppose that two people are talking to each other. One person says something along the lines, "There's this interesting paper I just read by some people at Berkeley about user inter-

Personal and Ubiquitous Computing (2001) 5:78–81

faces", and goes on to describe the paper in more detail. Even though the speech recognition is not perfect, enough of what was said was recognised. The speaker's CIA uses what was said, combined with the history of documents recently seen, to find potential matches, making the referenced paper easy to find and access if desired.

3. Information Retrieval

What we have described above is a process-oriented view, one that describes *how* the information is being retrieved. Another way of thinking about it is by the *relationship* of the information to the user. The information being retrieved can be thought of as 1) information a person would have searched for manually; 2) related information the person already knows; 3) serendipitous information the person didn't know; or 4) completely unrelated information (Fig. 1). Our goal is to maximise the first type, information that would have been searched for manually. From a systems standpoint, one can think of this as pre-fetching, streamlining what the person would have done anyway.

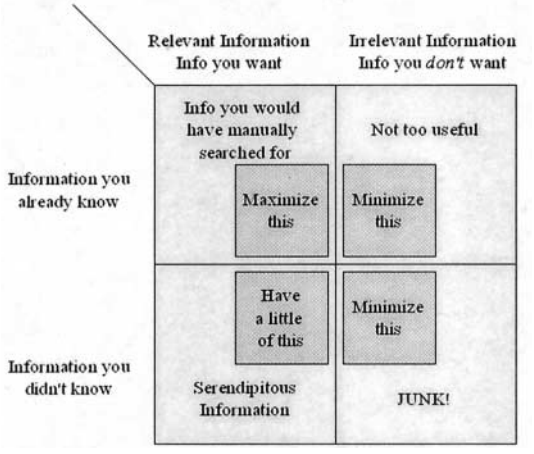

Fig. 1. CIA Information Retrieval Space.

However, getting the information is only part of the problem. Just as important is how to present the information in such a manner to support the task, without overly distracting the users. For example, a shared display of constantly updating results would simply be too disruptive in a meeting. A system that constantly beeps and presents marginally useful information would also be too distracting. The key is to balance the utility of the information with the disruption it would cause the users, an approach suggested by both Horvitz [1] and Maes [2].

4. Low Fidelity Prototype

Before implementing a system, we decided to run a low-fidelity prototype in a meeting situation to explore the domain and to test out some ideas. An audio recording was made of a weekly meeting. After the meeting, one of the authors did searches based on what was said. All of the

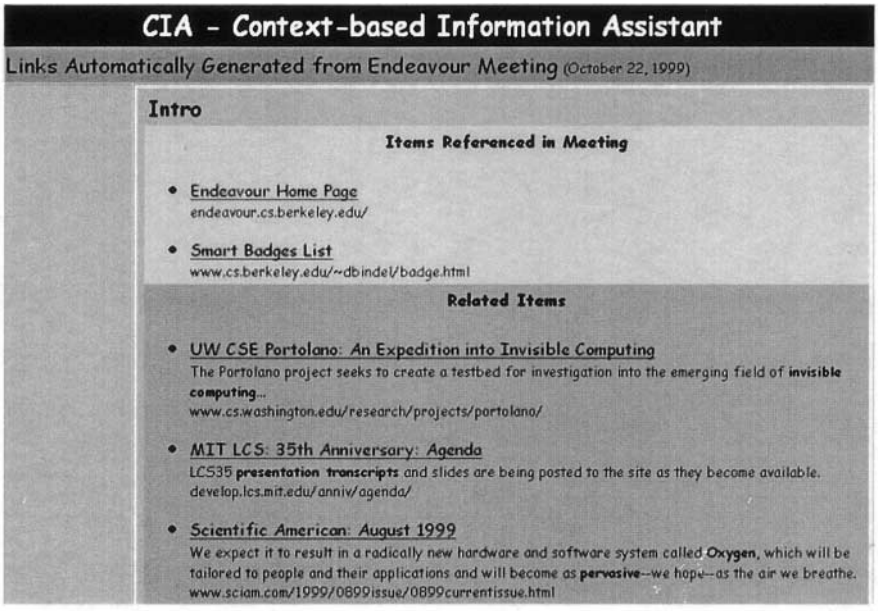

Fig. 2. Low-fidelity prototype of search results from meeting.

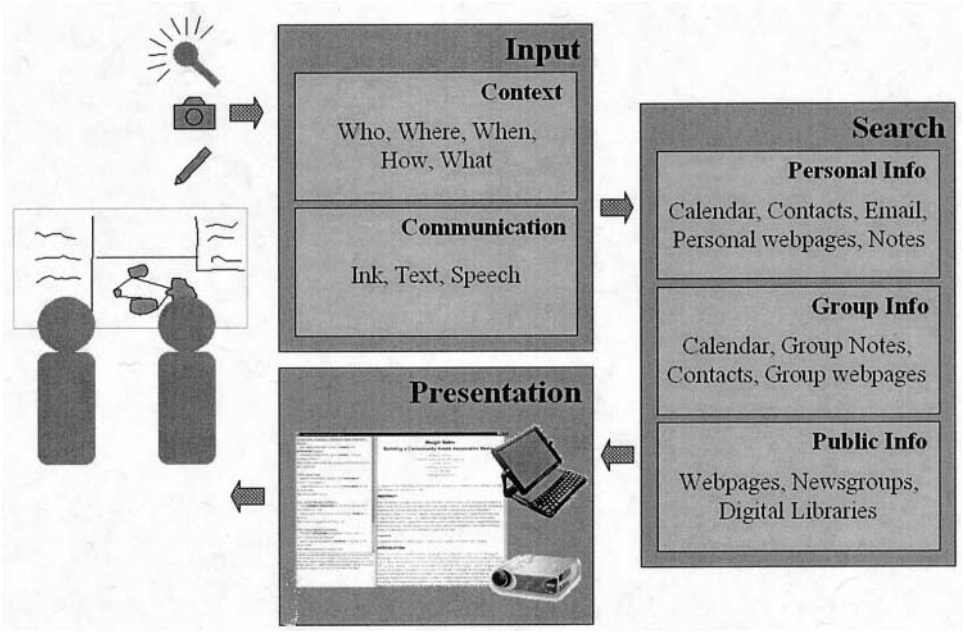

Fig. 3. Design space of CIA.

results were assembled into a web page, organised chronologically and by general topic (see Fig. 2). In each topic, the results were grouped by items explicitly referenced during the meeting, and items related to the discussion but never explicitly mentioned.

Once the results were organised, the meeting participants were asked to look over the results and fill out a short survey, judging the usefulness of the results as well as the organisation scheme. The general results were that people liked the concept a lot, but wanted more useful results, as well as more sophisticated ways of organising and filtering the results. Furthermore, people were interested in seeing if the system would be useful during a meeting real-time. One serious concern was control of the system: people should be able to turn it on and off when desired.

Taking all of what we learned from the scenarios and from the low-fidelity prototype, we decided to take a look at the design space (Fig. 3). As a first cut, we have broken the design space into three main components: Input, Search and Presentation. As stated earlier, the input is gathered by capturing the user's current context and communication. This information is used to search on various collections of information, some personal, some belonging to a group, and some publicly available. The results, once filtered, are measured for predicted utility and disruption to the user. If the information is useful enough, then the system calls the user's attention to it. Otherwise, it is just added to the list of marginally useful items, which the user can easily access if desired.

5. Speech-based Implementation

As a first step towards building CIA, we developed a prototype that takes speech input, processes it through a speech recogniser, and then does web searches based on keywords spotted in the recognised speech. It can be currently thought of as a speech-based interface for web search engines. We are presently in the process of improving the recognised speech, as well as expanding the search to other kinds of information, such as digital libraries.

We are also in the process of investigating several strategies to minimise attention to the agent in a real-time meeting situation. First, we believe that peripheral displays will be useful, i.e. using secondary monitors and projectors off to the side to display the results. Second, we believe that periodic updates will be more useful than continuous updates, so that people will not have to read constantly changing information. Third, we believe that pre-processing the results to extract the most important headers and text can significantly reduce the amount of reading

needed. In addition, there are intriguing directions to explore for asynchronous interaction, such as receiving an email from the agent after a meeting.

6. Related Work

In several respects, the CIA as envisioned is similar to Remembrance Agents [3], but moves the focus away from keyboard input and from wearable computers towards natural modes of communication. The CIA is also related to XLibris [4], a pen-based portable document reader specifically designed for reading electronic documents. One notable feature in XLibris is implicit linking: highlighting phrases in one document would cause the system to search locally for related documents. Any links found would be presented as a small document icon in the margin next to the highlighted text. Thus, the user never explicitly searches: instead documents are found opportunistically. The key observation is that useful information can be found based on activities one is already doing. The CIA also has a strong relationship with meeting capture systems, such as Classroom 2000 [5] and the data salvaging tools at PARC [6]. A CIA can be thought of as being built on top of these kinds of systems.

7. Summary

The CIA is a system we are working on that leverages context and human-to-human communication, using this knowledge to search for the right information at the right time. We are currently in the process of developing this system for a group meeting environment, and are investigating methods for expanding the collections of information the CIA searches, approaches for improving the presentation of the search results, as well as strategies for minimising the amount of attention the agent brings to itself.

References

1. Horvitz E. Principles of mixed-initiative user interfaces. In: Proceedings of Human Factors in Computing Systems: CHI'99. ACM Press, Pittsburgh, PA, 1999; 159–166
2. Maes P. Agents that reduce work and information overload. Communications of the ACM 1994; 37: 31–40
3. Rhodes B, Starner T. The Remembrance Agent: a continuously running automated information retrieval system. In: Proceedings of The First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology (PAAM '96). London, UK, 1996; 487–495
4. Schilit BN, Golovchinksy G, Price MN. Beyond paper: supporting active reading with free form digital ink annotations. In: Proceedings of Human Factors in Computing Systems: CHI'98. ACM Press, Los Angeles, CA, 1998; 249–256
5. Abowd GD, Atkeson CG, Brotherton J, et al. Investigating the capture, integration and access problem of ubiquitous computing in an educational setting. In: Proceedings of Human Factors in Computing Systems: CHI '98. ACM Press, Los Angeles, CA, 1998; 440–447
6. Moran TP, Palen L, Harrison S, et al. "I'll get that off the audio": a case study of salvaging multimedia meeting records. In: Proceedings of Human Factors in Computing Systems: CHI'97. ACM Press, Atlanta, GA, 1997; 202–209

Correspondence to: J. I. Hong, Group for User Interface Research, Computer Science Division, University of California, Berkeley, Berkeley, CA 94720-1776 USA. Email: {jasonh, landay}@cs.berkeley.edu