

A Context / Communication Information Agent

Jason Hong, James Landay
Group for User Interface Research
University of California at Berkeley
{jasonh, landay}@cs.berkeley.edu

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).

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 could be used, suppose that earlier in the day, Francis scribbled down a grocery list. Later, when passing by the grocery store he usually goes to, his PDA beeps, reminding him to buy some food. As he enters, his PDA fetches his handwritten notes for him. As another example, suppose that a person has a weekly meeting to go to, stored as a recurring weekly event on his calendar. When the time for the next meeting takes place, the system could begin retrieving notes, minutes, and action items from last week’s meeting, so that he doesn’t have to remember where he saved them.

By human-to-human communication, we mean using microphones, cameras, and other sensors to capture communication between people, such as text, ink, speech, and so on. As an example of how communication could be used to prefetch 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 interfaces”, and goes on to describe it more in detail. Using the information that was said, the system could begin searching for potential matches, so that the referenced paper, and possibly related papers, will be there if needed.

What we described above is a process-oriented view, that is it describes how the information is being retrieved. Another way of thinking about it is by the type of information being retrieved. The information being retrieved can be thought of as information a person would have searched for manually; related information the person already knows; serendipitous information the person didn’t already know; or completely unrelated and useless information. Our goal is to maximize the first type, information that would have been searched for manually.

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 display of constantly updating results would simply be too disruptive in a meeting.

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 results were assembled into a web page, organized chronologically and by general topic (see Figure 1). In each topic, the results were grouped by items explicitly referenced during the meeting, and items related to the discussion but never explicitly mentioned.

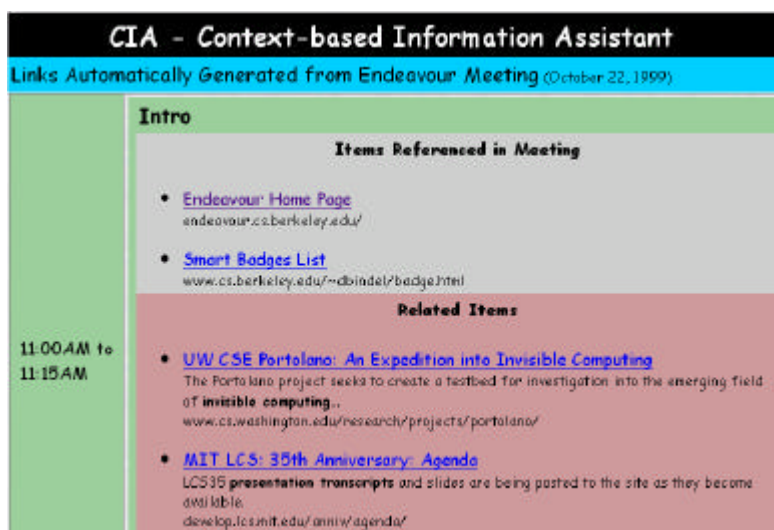


Figure 1 – Low-fidelity prototype of search results from meeting

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

Next, we built a prototype that takes speech input, processes it through a speech recognizer, and then does web searches based on keywords spotted in the recognized 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 recognized 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 minimize attention to the agent in a real-time meeting situation. First, we believe that peripheral displays will be useful, that is 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.

In several respects, the CIA as envisioned is similar to Remembrance Agents [1], but moves the focus away from keyboard input and from wearable computers. The CIA is also related to the XLibris system [2], 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: documents are instead 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 [3] and the data salvaging tools at PARC [4]. A CIA can be thought of as using the same infrastructure as these systems or built on top of these kinds of systems.

References

1. Rhodes, B., and Starner, T. The Remembrance Agent: A Continuously Running Automated Information Retrieval System. In the Proceedings of The First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology (PAAM '96), London, UK, April 1996. pp. 487-495. <http://rhodes.www.media.mit.edu/people/rhodes/research/Papers/remembrance.html>
2. Schilit, W.N., Golovchinsky, G., and Price, M. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations. In the Proceedings of CHI '98, 1998. ACM Press.
3. Gregory D. Abowd, Chris Atkeson, Ami Feinstein, Cindy Hmelo, Rob Kooper, Sue Long, Nitin "Nick" Sawhney and Mikiya Tani. Teaching and Learning as Multimedia Authoring: The Classroom 2000 Project. In the Proceedings of the ACM Multimedia'96 Conference, pp. 187-198, November 1996.
4. Moran, T., et al. "I'll Get That Off the Audio": A Case Study of Salvaging Multimedia Meeting Records. Proceedings of CHI'97, Atlanta, GA. ACM Press. 1997.