MeetingManager: A Collaborative Tool in the Intelligent Room

Alice Oh AOH@AI.MIT.EDU
Rattapoom Tuchinda PIPET@AI.MIT.EDU
Lin Wu LINWU@MIT.EDU
MIT Artificial Intelligence Laboratory, 200 Technology Square, Cambridge MA, 02139 USA

1. Introduction

The Intelligent Room Project (Coen, 1998) is an integral part of the MIT Oxygen Project, providing the overall software and hardware infrastructure for E21. The Room is equipped with several sensors and audiovisual hardware including projectors, electronic whiteboards, microphones and cameras, all of which are connected to the Metaglue System (Coen, et al. 1999). This setup makes the room a natural environment for smart collaboration. In this paper, we describe our MeetingManager system, a multi-user multimodal collaboration tool for planning, facilitating, and browsing structured meetings.

1.1 Towards Effective Meetings

Members of many organizations, large and small, spend a considerable amount of their precious time in meetings. However, it is rarely the case that everything that happens in a meeting results in a desired outcome, whether it be shared knowledge of some issue or an action item someone has committed to. To make every meeting more effective, an accurate record of the meeting should be kept. Moreover, the records must allow easy retrieval and efficient browsing. The MeetingManager is a robust system that assists in running effective meetings by taking advantage of the capabilities of the Intelligent Room.

1.2 Other Systems

Computer meeting support has been a topic of active research in the Computer-Supported Cooperative Work (CSCW) community. However, their focus has been modifying the meeting process via use of technology, rather than developing systems to better support the meeting process. For example, in (Nunamaker, et al. 1991), all communication is done via typed input from meeting participants, each sitting in front of a workstation. While this may increase participation and improve the group decision making process, it has the disadvantage that the communication style is very different from the natural-occurring meetings, thereby forcing all participants to learn and adapt to the new meeting process. In (Streitz, et al. 1994), they take a step further and support different types of meetings, including traditional face-to-face meetings, but the meeting record is again based on the written material produced during the meeting.

A team at Xerox PARC has looked at capturing multimedia meeting records and buidling tools for *salvaging* (Moran, et al. 1997). This requires the salvager to go through the multimedia meeting record after the meeting and produce an annotated record. A similar approach is taken in (Davis, et al. 1999), but with personal digital assistants (PDAs) and capturing only written notes. These two systems offer flexibility of supporting unstructured meetings, the latter providing support for meetings held anytime, anywhere (assuming the participants have PDAs).

The Interactive Institute has been working on a system for multimodal meeting record, in which audio and video are processed and annotated automatically with rich information such as speaker identity, participants' emotions, etc. (Gross, et al. 2000). Their system uses sophisticated technologies such as large-vocabulary speaker-independent speech recognition, eye-gaze tracking, and dialogue processing. However, while the paper mentions that it captures the action items when they are reiterated at the end of the meeting, the system itself does not impose and support any structure to a meeting. This makes it more difficult to use the meeting record as a simple reminder of the meeting to the participants.

Our MeetingManager takes advantage of the natural structure of the meeting that many organizations practice for an efficient meeting process. In some sense, it is a combination of the meeting process-focused work of the CSCW community, and the free-form meeting support of the HCI community. The MeetingManager takes advantage of minimal structure for a meeting (an agenda with issues and commitments) to produce automatic, real-time annotated multimedia meeting record.

2. Components of MeetingManager

The MeetingManager has four components that work together: a planner, a facilitator, a summarizer, and a browser. Before the meeting takes place, the participants use the planner to make and store the agenda. During the meeting, the facilitator takes the stored agenda for real-time meeting assistance. After the meeting, a brief summary of the recorded multimedia is emailed to the participants. Finally, the meeting is stored in a database available for convenient browsing. In this section, we describe the four components in more detail.

2.1 Meeting Planner

Before the meeting, the planner assists in creating an agenda and storing it into a persistent database. Each pending meeting is tagged with its owner and the time and date of the meeting. An agenda is a Java object with a list of participants and a list of agenda items. Each of the agenda items is itself a Java object with the item description, duration (proposed length of discussion), priority (high, medium, low), and state (active, complete, incomplete, skip). The state of an agenda item allows it to be continued at a later time (e.g., the next scheduled meeting) without manually re-entering it into the agenda.

2.2 Meeting Facilitator

During the meeting, the facilitator displays the agenda, times each item according to its duration, and videotapes the discussions, which is then saved as a QuickTime video file. For each agenda item, the facilitator also records issues and commitments raised by the participants, along with the time stamps of those events. The issues and commitments are also displayed real-time alongside the agenda. The agenda items can be checked off when completed, or they can be skipped or reordered according to time constraints. The facilitator uses multimodal user interface including speech (IBM ViaVoice), typing, and electronic whiteboard (Mimio).

2.3 Meeting Summarizer

At the end of the meeting, the MeetingManager automatically sends a summary to each of the participants. The summary contains the agenda, the issues, and the commitments, along with the hyperlinks to the sections of the QuickTime video for those events. The participant can click on the links to play back the clips.

2.4 Meeting Browser

The agenda object, the issues, the commitments, and the QuickTime file are all wrapped into a meeting object and stored into a database, which can be queried later on using SQL. This allows a quick review of previous meetings as well as an efficient lookup of all relevant events. The browser, as well as the planner and the facilitator, supports natural language input using a tailored grammar.

3. Future Plans

The MeetingManager is a well-suited platform for carrying out a variety of multimodal human-computer interaction (HCI) experiments. We plan to add components for a more natural HCI including hand gesture recognition, speaker detection and tracking, and non-verbal discourse cue recognition. With more sophisticated computer vision algorithms and HCI tools, the MeetingManager will become an indispensable application of ubiquitous computing.

Acknowledgements

This work is done with collaboration from the members of the Intelligent Room Project, the Vision Interface Project, and the Design Rationale Group. Krzysztof Gajos, Stephen Peters and Andy Chang especially deserve credit for their hard work. Thanks to Dr. Howard Shrobe and Prof. Trevor Darrel for guiding the project with vision.

References

Coen, M. (1998). Design principles for intelligent environments. *The Fifteenth National Conference on Artificial Intelligence. (AAI98)*. Madison, Wisconsin.

Coen, M. et al. (1999). Meeting the computational needs of intelligent environments: The metaglue system. *MANSE 1999 International Conference on Managing Interactions in Smart Environments*.

Davis, R., et al. (1999). NotePals: Lightweight Note Sharing by the Group, for the Group. *Proceedings of Human Factors in Computer Systems: CHI 99*,

Gross, R., et al. (2000). Towards a Multimodal Meeting Record. *Proceedings of ICME 2000*.

Moran, T. P., et al. (1997). I'll Get That Off the Audio: A Case Study of Salvaging Multimedia Meeting Records. *Proceedings of Human Factors in Computing Systems, CHI97,* Atlanta, GA, March 22-27, 1997. pp. 202-209.

Nunamaker, Jr., J.F., Briggs, R.O., Mittleman, D. (1991). Electronic Meetings to Support Group Work. *Communications of the ACM*, 34(7)1991, pp.40-61.

Streitz, N., et al. (1994). DOLPHIN: Integrated Meeting Support across LiveBoards, Local and Remote Desktop Environments. *Proceedings of the Conference on Computer-Supported Cooperative Work, CSCW'94*.