The Necessity of a Meeting Recording and Playback System, and the Benefit of Topic–Level Annotations to Meeting Browsing

Satanjeev Banerjee, Carolyn Rose, and Alexander I. Rudnicky
{banerjee+, cprose, air+}@cs.cmu.edu

Language Technologies Institute
Carnegie Mellon University

Abstract. In modern organizations, meetings are the forum set aside for sharing and producing information that is crucial to the conduct of business. Much work in the area of Computer Supported Cooperative Work (CSCW) has targeted the problem of supporting meetings between collaborators who are non-collocated, enabling meetings to transcend boundaries of space. In this paper, we explore the beginnings of a proposed solution for allowing meetings to transcend time as well. The need for such a solution is motivated by a user survey in which busy professionals are questioned about meetings they have either missed or forgotten the important details about after the fact. Our proposed solution allows these professionals to transcend time in a sense by revisiting a recorded meeting that has been structured for quick retrieval of sought information. Such a solution supports complete recovery of prior discussions, allowing needed information to be retrieved quickly, and thus potentially facilitating the effective continuation of discussions from the past. While prior work on video recorded meetings has explored the use of annotations attached to strategic points in the video record as retrieval aid, we explore the potential to develop fully automatic annotations on meeting recordings based on the discourse structure of the ensuing dialog. We evaluate this proposed solution with a formal user study in which we measure the impact on the proposed structural annotations on retrieval of information. The results of the study show that participants took significantly less time to retrieve the answers when they had access to discourse structure based annotation than in a control condition in which they had access only to unannotated video recordings ($p < 0.003$, effect size 0.94 standard deviations).

1 Introduction

Meetings in modern organizations play a crucial role as the forum where information is shared, alternatives are discussed and decisions are made. A lot of research has therefore focused on facilitating the conduct of meetings when all the participants are not collocated (e.g.: [11], [12], [13], etc). A secondary role that meetings play is in the bringing together of a large amount of project–relevant information, and also being the source of new kinds of information (such

as decision, action items, etc). Some recent research has focused on capturing this information, both old and new, and making it available for later access by interested parties (e.g.: [1], [4], [6], [7], etc.) Such technology has several potential uses. First they can be of use to parties that were absent during the actual meeting. While such *non–contemporaneous* meeting participants may not be able to contribute to the meeting itself, they can at least benefit from the information produced at the meetings using such technology, and then be included as a more informed participant in resulting ongoing discussions. A second use of this technology is as an aid to participants' memories of past meetings. Participants forget details of meetings, or worse, have erroneous recollections of past meeting. Meeting capture and play back technology, if appropriately designed, can be used to efficiently retrieve details of past meetings. Thus this technology can be viewed as improving Organizational Memory [8], which has been shown to improve users' productivity [9].

Our first goal in this paper is to explore the problem of reconstructing information discussed at a meeting. How often do busy professionals need to reconstruct details of past meetings? What kinds of documents do they typically have access to? Are those documents sufficient? What kinds of information are they typically seeking from past meetings? How much time does it take to do the reconstruction? To gain an understanding of these issues, we have run an interview–based survey with 12 faculty members at Carnegie Mellon University. We have chosen this user population since university faculty typify professionals who's lives are dominated by meetings. Interviewees were asked to narrate specific instances of situations when they were trying to catch up on a meeting that they had missed, or were trying to reconstruct forgotten details of a meeting they had attended in the past.

Our second aim is to guage how helpful topic–level annotations are to meeting browsing. The Carnegie Mellon Meeting Recorder [1] can be used to create rich multi–modal records of meetings, while the MockBrow can be used to manually annotate a meeting record, and also play back both the meeting record along with any associated annotations. Both these tools are being developed as a part of the CALO (Cognitive Agent that Learns and Organizes) project [16]. The ultimate goal is for the discourse structure based annotation that we evaluate the utility of in this paper to be applied to recordings of meetings automatically by making use of easily detectable cues such as shifts in word frequency distributions that are both indicative of topic shifts.

For the purposes of this study, we use the MockBrow to annotate meeting records by labeling different portions of the meeting with their general "topic of discussion". We also mark different parts of the meeting with what the "state" of the meeting was (discussion / presentation / briefing). The meeting is not annotated with any other information, such as ontology based structural information as in [10]. We hypothesize that a human can retrieve information from a meeting faster if he is armed with topic annotations versus if he is not. In section 4 we present a user study aimed at quantifying the extent to which this hypothesis is true.

2 Related Work

Three previously published lines of research form the foundation for the investigations reported in this paper.

The Distributed Meetings (DM) project [4] involves a system that can be used to broadcast and record meetings, and also view pre–recorded meetings. Unlike our system, the DM does not attempt to automatically identify the structure of the meeting. The author of this project evaluated the DM by conducting a user study involving real meetings between real participants at Microsoft Research. At each meeting, a person was asked to remain absent, and to later come in and view the meeting recording using the DM viewing software, and then fill out a questionnaire. The questionnaire data did not present a strong case for the desirability of the solution although it did provide evidence that users were satisfied with the information they received from it. What the paper lacks that we attempt to provide is insight into the specific types of information needs people have regarding missed or forgotten meetings.

Another similar project is the WorkspaceNavigator [7] that attempts to capture many different sources of digital information as a meeting proceeds inside a "smart room". Recording involves taking regular snapshots of different computer displays, the meeting room itself, filenames of open files and URLs of visited webpages from participant laptops, etc. Users are allowed to label the snapshots, or just mark snapshots as being important as they are being recorded. Two qualitative user studies were conducted to provide a detailed view of patterns of actual use of the technology. The paper provides convincing evidence that users were able to index and retrieve portions of meetings when needed. What this previous paper lacks that we offer in this paper is a controlled experiment in which the magnitude of the impact of a general class of annotations on the speed of information retrieval from meeting recordings is precisely measured.

Closely related to our project is the survey [5] of potential users of a meeting browser, conducted as a part of the IM2.MDM (Interactive Multimodal Information Management, Multimodal Dialogue Management) project. The goal of this survey was to elicit a set of questions that users may ask of an intelligent meeting browser. Participants were asked to imagine themselves in one (or more) of four roles – a manager tracking employee performance, a manager tracking project progress, an employee who has missed a meeting, and a new employee – and to then think of all the questions they expect to ask about the meeting or set of meetings that the meeting browser may have access to. While this survey provides some broad insights, it differs from ours in both its goals and its methodologies, especially in that it does not adhere to strong HCI methodology for survey research. One of the goals of our survey is to assess how *useful* a meeting browser would be, how urgently its need is currently felt by busy professionals, and in what range of their actual situations they could potentially benefit from the use of a meeting browser. In contrast, the survey reported in [5] makes the implicit assumption that if busy professionals had access to an intelligent meeting browser, they would indeed use it! In our survey interviewees were asked to recall recent instances of actual situations. The resulting analysis of the

interviews is therefore grounded in real experiences as opposed to potentially erroneous generalizations. The survey questionnaire in [5] asked participants to imagine themselves in a situation they have never been in before, namely, in possession of a system using which they could "ask questions about the actual content of the meeting". Thus, it is not clear how many of these questions would indeed be asked by users of a future meeting browser in actual use.

3 User Survey

3.1 Goals

Efforts towards creating a meeting recording and play–back system can clearly drive research on a large number of fronts including speech recognition, spoken language understanding, vision–based gesture recognition, multi–modal information integration, etcetera. In this paper however we are interested in exploring the *need* for such a meeting browsing application. Intuitively, such an application would be useful to busy professionals who need to catch up on missed meetings or recall forgotten details of meetings they have attended in the past. To understand how professionals currently perform these tasks we have conducted an interview based survey. Specifically, the survey was conducted to find answers to the following questions:

- How often do busy professionals miss important meetings that they need to catch up on?
- How often do users need to reconstruct forgotten details of meetings they did attend?
- What kind of information/documents do they typically have access to in each of the above two cases?
- What kind of information do they typically seek?
- What processes do users currently employ in obtaining this information?
- How effective are these processes in terms of accurately retrieving the desired information?
- How costly are these processes in terms of time/energy spent on retrieving the information?

3.2 Survey Methodology

Our survey was based on face–to–face interviews conducted with 12 faculty members at Carnegie Mellon University. We chose faculty members since they attend many meetings as a part of their daily routine, and would be the likely targets of a meeting recording and playback application. Since not all missed meetings are important enough to bother catching up on, we defined a meeting as *important* if the interviewer felt he would indeed make an attempt to find out about it if he missed it. Interviewees were asked to describe instances of two kinds of situations: situations when they had missed important meetings, and those in which they were trying to recollect details of a meeting they had attended in

the past. For each instance, interviewees were asked to name and describe the meeting artifacts they had access to, what specific pieces of information they were seeking about the meeting, whether and how they found the information, whether they were satisfied with the information they did find, etc. To avoid bias, interviewees were not informed about the reasons for this interview until the very end of the interview. To ground the interview in real experiences, interviewees were strongly and repeatedly encouraged to avoid replying in potentially erroneous generalities, and instead were asked to recall specific situations from their experiences in answering questions.

3.3 Analysis of Non-Missed Meetings

The 12 interviewees reported a total of 19 instances of situations when they were attempting to recall details of a meeting they had attended in the past (1 interviewee reported no such instances, 3 interviewees reported 1 each, and the rest reported 2 each). Interviewees were asked to report both when they were attempting to recall details of a past meeting, as well as when the meeting took place (which was normally within the past few months).

Information Sought from Meeting: Interviewees were asked to specify the information they were attempting to reconstruct about the meeting; table 1 lists the frequencies of the various categories of information sought across all the instances of non–missed meetings reported by the interviewees. Note that interviewees were not shown the categories listed in the table, but were simply asked to recall all the pieces of information they were seeking about the meeting. These answers were later manually clustered into the groups in table 1. For example, the category *Specifics of the discussion on a topic* include questions like "What was the name of the algorithm we discussed?".

While interviewees were not specifically asked to explain why they needed the information they were seeking, for several of the 7 instances of the category *What the decision was regarding a particular topic* interviewees spontaneously mentioned that the reason they were attempting to recall the decision was not because they thought they had forgotten the detail, but because their recollection of the detail differed from that of another co-participant of that meeting. We believe that this phenomenon is an important motivation for meeting recording and play–back technology.

Reconstructing from Available Documents: The interviewee was asked to list the documents he had access to while he was attempting to make the reconstruction; table 2 lists the documents. In 14 of the 19 meetings, the users had access to notes taken at the meeting, typically the notes they had taken during the meeting. In the remaining instances, interviewees had not taken notes at the meeting, and further did not have access to notes taken by any other meeting participants. Interviewees were also asked to rate on a scale of 0 to 5 whether the piece of information they sought about the meeting was satisfactorily answered

Table 1. Information Sought from Meeting.

Information sought	# meetings
Specifics of the discussion on a topic	11
What the decision was regarding a particular topic	7
What task someone else was assigned	4
Who made a particular decision	2
What the participants' reactions were to a particular topic	1
What the future plan is	1

by the meeting documents they had at their disposal, where 0 implied their question remained unanswered, and 5 implied they were completely satisfied with the answer they got. The average rating was 3.0 (std. dev.: 1.7).

Table 2. Documents the Interviewee Had Access To.

Documents interview had access to	# meetings
Notes	14
Nothing	2
Minutes	1
PowerPoint Slides	1
Excel Sheet	1
Project proposal document	1
Whiteboard content	1
Email	1

Additional Steps Taken to Find Information: Interviewees were asked what additional steps (besides perusing the meeting documents) they took to find the information they needed from the meetings. In 8 cases the interviewees asked someone in a face–to–face conversation. This was particularly the case when the question was about a specific detail about the meeting. In 5 cases interviewees reported that they reconstructed from memory, in consultation with a meeting co–participant (note that this is not the same as simple *asking* someone else about a detail). Finally interviewees were asked to rate on a scale of 0 to 5 their quality of reconstruction of the information they were seeking, after they took the additional steps, where 0 implied they could not do any reconstruction at all. The average rating was 4.0 (std. dev.: 0.6) – this was significantly higher than the satisfaction before perusing the meeting documents ($p < 0.0005$). 5 interviewees stated that the additional steps took less than 15 minutes, 7 said between 15 minutes and an hour, while for 2 interviewees, the additional steps took more than an hour.

Summary and Conclusions:

- Interviewees mostly sought very minutely detailed pieces of information from the meetings they had attended in the past.
- Very often the interviewees had access to notes that they could consult.
- Interviewees sometimes felt satisfied with the information they were able to retrieve from available documents.
- When the documents did not suffice, interviewees spoke to co–participants, or took other additional steps, which took up to an hour of time. At the end of these steps, interviewees largely felt that their information needs had been satisfied. Nevertheless, this does not guarantee that the information that they received was accurate since we have already established that meeting co-participants may have different recollections of what was discussed at a meeting.

3.4 Analysis of Missed Meetings

The 12 interviewees reported a total of 22 instances of meetings they had missed in the past that they needed to catch up on. 9 interviewees reported 2 instances each, while 1 interviewee reported 3, 1 2, and 1 none (that is, one interviewee could not recall any specific instance of an important meeting that he had missed and later attempted to catch up on). 2 of these 22 missed meetings had occurred in the week prior to the interview, while 10 had occurred within the preceding month. Thus, on average interviewees missed one important meeting a month. Of the remaining instances, 9 had occurred within six months prior to the interviews, and 1 between six months to a year before the interview. Note that based on the frequency of missed important meetings reported within the month prior to the survey, it would not be unrealistic to estimate that the population used in the survey missed on average almost 1 important meeting per month.

Understanding of Expected Meeting Content Prior to Meeting: A person's overall understanding of the information discussed at a meeting is likely affected by his prior knowledge and expectations about the meeting before it takes place. Of the 22 reported instances of missed meetings, in 2 cases the interviewee did not receive any notification about the meeting (such as an email announcing the meeting or inviting him to it). In one of these cases the meeting had already taken place by the time the interviewee received the notification, while in the other case the meeting was a regularly scheduled one and notifications weren't usually sent out. In 12 of the remaining cases the interviewee received an agenda and/or a description of what would be discussed, while in the remaining cases he received a notification email announcing the meeting. Each interviewee was asked to rate on a scale of 0 to 5 how well he felt he knew the contents of the meeting would be, where 0 implied he had no idea what the contents would be, and 5 meant he knew exactly what would be discussed. The average rating was 3.5 (std. dev.: 1.3).

Information Sought from Missed Meeting: For each instance of reported missed meeting, interviewees were asked why they wanted to catch up on the meeting. In particular, they were asked to list all the pieces of information they needed from each missed meetings. Table 3 presents the frequency of each category of information sought by the interviewees across all the instances of missed meetings. Thus in 10 of the 22 missed meetings, the interviewee wished to find out what was discussed about a specific topic. As with non–missed meetings, participants were not provided with the categories in table 3; the categories were constructed based on their responses. [Note that in the last 2 interviews the interviewee was indeed provided the categories, but also encouraged to add information if they felt the categories were not enough. Neither interviewee felt the need to add to these categories].

Observe that the first 3 categories in table 3 together make up the majority of categories of information sought about missed meetings; note also that these three categories are related in that they are all concerned about seeking information about a particular topic of interest. This suggests that when a person misses a meeting, he is often more interested in catching up with the discussions regarding a specific topic of interest rather than the entire meeting. Thus perhaps an automated topic detection and segmentation mechanism that lets the viewer of a recorded meeting focus only on the topic he is interested in will be well received. Surprisingly, there were only a few cases in which interviewees wanted to know about progress related information like action items, announcements, etc. We conjecture that this is because most of the meetings reported were research discussion–oriented where the discussion details were perhaps more important to the interviewee than action items.

Table 3. Information Sought from Missed Meetings.

Information sought	# meetings
What was discussed about a particular topic	10
What decisions were made about a particular topic	7
Whether a particular topic was discussed	5
Whether I was assigned a new task	4
Whether a particular decision was made	3
What decisions were made	2
If there were any new issues/announcements	2
Reasons for a decision	1
What the participants' reactions were to a particular topic	1
The backgrounds of the other participants	1

Understanding of Meeting Content After the Meeting: To understand what kind of information about the meeting the interviewees could access with-

out having to make an effort at locating the information, we asked the interviewee to list the documents he received from the meeting after the meeting took place, *without him prompting for them.* Note that by "document" we included *any* piece of information that the interviewee may have received without prompting, including oral intimation from other participants of the meeting, emailed documents, etc. The aim in asking this question was to understand what kind of documents are routinely sent around – one can presume that (at least a large subset) of these documents would be available even when participants have access to a meeting browsing system. Table 4 lists the documents reportedly received by the interviewees. Observe that in more than half the meetings, no document was received at all. Since these meetings are important enough to the interviewee that he wants to catch up on them, not receiving any information from the meeting implies that the interviewee is forced to either actively search for information about the meeting, or give up and not learn anything about the meeting at all.

Interviewees were further asked to rate on a scale of 0 to 5 how well the documents they received (if any) answered their question(s) about the meeting, where 0 meant they either did not receive any documents, or that the documents they received did not answer their questions at all. The average rating was a very low 1.7 (std. dev.: 1.9). This low number is partly explained by the fact that to a large extent interviewees received nothing from the meetings. In cases where the interviewee did receive notes etc from the meeting, they were often not sufficiently detailed to answer their questions regarding the meeting.

Table 4. Documents Received after the Meeting.

Post–meeting document received	# meetings
Nothing	12
Notes	7
Minutes	3
Email from meeting participant (not official notes)	2
Document containing draft of a proposal	1

Additional Steps Taken to Find Information from Missed Meeting Interviewees were asked what additional steps, if any, they took to find answers to their questions regarding the meetings. Table 5 shows the steps taken. Consistent with our findings about non-missed meetings, in connection with 15 meetings interviewees either asked someone face to face about the meeting, or emailed someone.

When asked how long these steps took, in 11 instances the interviewee said it took less then 15 minutes, in 5 cases between 15 minutes and an hour, and in 1 case more than an hour (this information was not collected for 5 instances). These

Table 5. Additional Steps Taken.

Additional step	# meetings
Asked someone face-to-face	9
Emailed someone	6
No additional steps	5
Caught up at next meeting	3
Looked up information on the Internet	1

times were self reported by the interviewees and are rough estimates only: the interviewees often reported having discussed other issues with their interlocutors while catching up on the missed meeting.

Finally the interviewees were asked to rate on a scale of 0 to 5 how much they believed their information need was met after taking the additional step, where 0 meant their information need was not met at all. The average rating in this case was 3.4 (std. dev.: 1.3).

Summary and Conclusions:

- Interviewees were more interested in catching up on discussion regarding specific topics rather than the entire meeting.
- Very often no documents were received, even though the meetings were important.
- Typically interviewees had a low level of understanding of the meeting from the documents received.
- Most interviewees attempted to answer their questions regarding the meeting by asking or emailing a co-participant. This extra effort took around 15 minutes.
- Even after taking additional steps to find information about the meeting, the interviewees' levels of understanding about the meeting were felt to be far from perfect.
- Based on the fact that information is often sorely lacking about a missed meeting, we conclude that a meeting recording and playback system would be useful for busy professionals to catch up on missed meetings. Further, if the meeting recording is segmented into discussion topics, users can focus on only their particular topics of interest, thus increasing their efficiency of extracting information from the meeting.

4 Meeting Browsing using Topic Annotations

In the second part of this paper we report on our investigation into the effect of meeting annotation on the time it takes for a user to retrieve information from a meeting.

4.1 Meeting Annotation

We are interested in automatically detecting the structure of a meeting. For the purposes of this paper, we define meeting structure on two levels: A coarse level consisting of meeting states and participant roles, and a finer level consisting of discussion topics.

We build upon a previously published ontology of meeting states and participant roles [2] based on extended observation of natural meetings between human beings. In this topology, there are three kinds of meetings states, as follows: Discussion state which is described as being a state in which a group of two or more meeting participants are involved in quick back and forth of discussion on a topic; Presentation state which is described as being a state in which one particular meeting participant is presenting information to the others in a formal setting; and Briefing state which is described as being a state where one participant is giving information to one or more meeting participants, but without involving either the formality of the presentation state, or the quick back and forth of the discussion state. Within each meeting state, the possible roles of the meeting participants are defined as follows: within the discussion state participants may take the role of discussion participants or of observer; within the presentation state presenter or observer and within the briefing state information-provider, information consumer and observer.

Discussion topic regions are defined as all the times of the meeting that are devoted to discussing a certain topic. Although "topic" itself can be defined on various levels of granularity, in general we are interested in broad high level topics such as those that typically form different agenda items. For example "buying a printer" may be considered a topic. In well structured meetings it is possible to split a meeting into a series of time segments, each segment containing all the discussion regarding a particular topic. Research has been done on automatically finding topics both in written texts [15] and in broadcast news [14]. While we are currently applying the text topic detection techniques described in [15] to media recorded at meetings, in this paper we use only manually annotated meetings to ensure high quality topic boundaries.

4.2 Brief Description of the MockBrow

We have implemented a first version of a multimodal meeting annotator and browser tool called the MockBrow. This tool is meant to be a generic playback of all information sources recorded at a meeting. The only requirement is that all the media be accurately time stamped. The software includes an interface to create annotations. Annotations are defined as intervals of time within the meeting that are associated with some pieces of information. For example, a human annotator can select a time interval within the meeting and label it as a "presentation". Similarly, intervals in the meeting may be labeled as belonging to discussion on a specific topic, like "buying a printer". The MockBrow has been designed to be extensible both in terms of the media it can play back and in the annotations it can support.

4.3 User Study Goals and Methodology

In order to assess the value of developing technology to automatically annotate meeting recordings with discourse structure based annotations in order to facilitate extracting information from recorded meetings, we designed a within subjects user study.

Materials: We created an audio-video record of two ten-minute-long meetings, each involving three participants. Next we manually annotated each meeting with meeting states, participant roles and discussion topics using MockBrow. The same meeting could then be viewed either with the annotations or without. Finally we prepared for each meeting a set of five questions. (To avoid biasing the questions, the annotations were not consulted while constructing the question set). The questions were objective in nature, each with a single correct answer – this made the users' answers easier to grade as correct or not. The questions were based on the types of questions collected as part of our survey research.

Participants: 16 Carnegie Mellon graduate students participated in the within subjects experiment. Note that while the survey was conducted on faculty members who would be the most likely target population for a meeting recording and playback application, the career status of the participants is not likely to affect the speed with which they are able to retrieve answers to questions that are provided for them from recorded meetings. Thus, for the purpose of this small study, it was sufficient to use students as participants. Nevertheless, our plan is in the future to invite busy faculty members to participate in a long term longitudinal study of meeting recording and playback.

Experimental Manipulation: Each participant was asked to answer the questions for each of the two meetings by viewing the meeting's video using MockBrow while searching for the answers. In order to control for ordering effects, subjects were randomly assigned to 4 configurations in which half of the subjects viewed the first meeting and then the second meeting whereas the other half of the subjects viewed the meetings in the opposite order. Furthermore, half of the subjects viewed the annotated version of meeting one and the unannotated version of meeting two, whereas the other half of the subjects viewed the annotated version of meeting two and the unannotated version of meeting one.

In all cases, participants were encouraged to answer the questions as fast as possible, and their time to completion of each meeting viewing and question answering was recorded.

4.4 User Study Results and Analysis

The timing data collected in the experimental manipulation was sufficient for comparing average speed of answering questions with and without discourse structure based annotations. The control group (the group that did not have access to the annotations) took an average of 10.0 minutes (std. dev. = 2.6) to answer the given questions, while the experimental group took 7.5 minutes (std. dev. = 1.4). A two-tailed Student's T-Test assuming unequal variance shows that this difference in the means is significant with $p < 0.003$. Further the effect

size (using the standard deviation of the control group as the denominator) is 0.94. This establishes that the difference in time taken to answer the questions when the participants could view the annotations versus when they could not is a reliable difference. Specifically, the annotations allowed participants to take on average 2.5 minutes less time to retrieve the answers to 5 questions in meetings that were 10 minutes long, when compared to those who had no access to the annotations.

5 Conclusions and Future Work

In the first part of this paper, we have reported on a user survey aimed at understanding how busy professionals such as faculty members deal with situations when they are attempting to catch up on missed meetings, or attempting to recall details of meetings they have attended in the past. One important finding was that the busy professions participating in our survey research missed on average 1 important meeting per month. Furthermore, they frequently discovered that their recollection of a discussion at a meeting was not consistent with another group member's recollection. The most frequent recourse when faced with a perceived need to recover meeting was to talk to a group member who was at the meeting. However, even in the case where people felt satisfied with the information received from a group member, it is not clear to what extent the information they receive based on another's recollection is accurate. Thus, the survey research provides some support for the usefulness of a meeting browser. It also provides an ontology of question types that represent the types of information typically sought by our target user population.

In the second part of this paper we have reported on a within-subjects user study performed to quantify the impact that discourse structure based annotations have on the time it takes users to retrieve the answers to focused questions from recorded meetings. We have shown that in our experiment, participants on average took 2.5 minutes less to find answers when given the annotations than when not, and that this is a highly significant difference ($p < 0.003$ using Students' two tailed T-Test, assuming unequal variance). This encouraging results is only a first step towards understanding the impact of discourse based annotations on retrieval of information from recorded meetings. We plan to perform a larger study with longer meetings, and with two populations of participants – those who have been in the meeting and need to recall details about the meeting, and those that have missed the meeting and need to find out about the meeting.

References

1. Banerjee, S., Cohen, J., Quisel, T., Chan, A., Patodia, Y., Al-Bawab, Z., Zhang, J., Black, A., Stern, R., Rosenfeld, R., Rudnicky, A.I., Rybski, P., Veloso, M.: Creating Multi-Modal, User-Centric Records of Meetings with the Carnegie Mellon Meeting Recorder Architecture. In: Proceedings of the ICASSP 2004 Meeting Recognition Workshop, May 17, 2004, Montreal, Canada.

2. Banerjee, S., Rudnicky, A.I.: Using Simple Speech-Based Features to Detect the State of a Meeting and the Roles of the Meeting Participants. In: Proceedings of the 8th International Conference on Spoken Language Processing (Interspeech 2004 - ICSLP), October 4-8, 2004, Jeju Island, Korea.
3. Rybski, P., Banerjee, S., Torre, F., Vallespi, C., Rudnicky, A.I., Veloso, M.: Segmentation and Classification of Meetings using Multiple Information Streams. In: Proceedings of the Sixth International Conference on Multimodal Interfaces, October 14th-15th, 2004, State College, Pennsylvania.
4. Cutler, Rui, Gupta, Cadiz, Tashev, He, Colburn, Zhang, Liu, Silverberg: Distributed Meetings: A Meeting Capture and Broadcasting System. In: Proceedings of the ACM Multimedia Conference, 2002.
5. Lisowska, A.: Multimodal Interface Design for the Multimodal Meeting Domain: Preliminary Indications from a Query Analysis Study. In: Technical Report IM2.MDM-11, November 2003.
6. Chiu, P., Kapuskar, A., Wilcox, L., Reitmeier, S.: Meeting Capture in a Media Enriched Conference Room. In: Proceedings of the Second International Workshop on Cooperative Buildings, Integrating Information, Organization, and Architecture, Pages 79-88, 1999.
7. Ionescu, A., Stone, M., Winograd, T.: WorkspaceNavigator: Tools for Capture, Recall and Reuse using Spatial Cues in an Interactive Workspace. In: Stanford Computer Science Technical Report 2002-04 (2002).
8. Stein, E. W., Zwass, V.: Actualizing Organizational Memory with Information Systems. In: Information Systems Research, Volume 6, Number 2, June 1995, pp. 127-137.
9. Jennex, M. E., Olfman, L.: Organizational Memory/Knowledge Effects on Productivity, a Longitudinal Study. In: 35th Annual Hawaii International Conference on System Sciences (HICSS'02) - Volume 4, January 07-10, 2002, Big Island, Hawaii.
10. Bachler, M. S., Buckingham Shum, S. J., De Roure, D. C., Michaelides, D. T., Page, K. R.: Ontological Mediation of Meeting Structure: Argumentation, Annotation, and Navigation. In: Proceedings of 1st International Workshop on Hypermedia and the Semantic Web (HTSW2003), Nottingham, UK, 2003.
11. Ackerman, M. S., Starr, Brian, Hindus, Debby, Mainwaring, Scott: Hanging on the 'wire': a field study of an audio–only media space. In: ACM Transactions on Computer–Human Interaction (TOCHI), 4, 1 (March 1997), 39–66.
12. Ahuja, S. R., J. Ensor, Robert, Horn, David N.: The Rapport multimedia conferencing system. In: ACM SIGOIS and IEEECS TC–OA Conferenceon Office information systems, ACM SIGOIS Bulletin, 9, 2-3 (April 1988).
13. Cook, Peter, Ellis, Clarence, Graf, Mike, Rein, Gail, Smith: Project Nick: Meetings augmentation and analysis. In: ACM Transactions on Information Systems (TOIS), 5, 2 (April 1987), 132-146.
14. Allan, J., Carbonell, J., Doddington, G., Yamron, J., Yang, Y.: Topic detection and tracking pilot study: Final report. In: Proceedings of the DARPA Broadcast News Transcription and Understanding Workshop, 1998.
15. Hearst, M.: TextTiling: Segmenting Text into Multi–Paragraph Subtopic Passages. In: Computational Linguistics, 23(1), pp. 33–64, March 1997.
16. http://www.ai.sri.com/project/CALO.