

# Visualization of Discussions in Face-to-Face Meetings

**TSUCHIDA, Takahiro**

Graduate School of Information Science  
tsuchida@nagoya-u.jp

**OHIRA, Shigeki**

Information Technology Center  
ohira@nagoya-u.jp

**KIUCHI, Keisuke**

Graduate School of Information Science  
kiuchi@nagao.nuie.nagoya-u.ac.jp

**NAGAO, Katashi**

Graduate School of Information Science  
nagao@nuie.nagoya-u.ac.jp

Nagoya University  
Furo-cho, Chikusa-ku, Nagoya 464-8603, Japan  
011-81-52-789-5878

## Abstract

*We previously developed a technique, called “discussion mining,” for semi-automatically generating content from a discussion. The content includes text information, audio-visual media, and metadata. We have now developed two systems for visualizing this content. A “discussion visualizer” system visualizes the structure of the discussion in the current meeting, and a “discussion reminder” system enables the contents generated for previous meetings to be retrieved and browsed. These systems facilitate understanding of the contents of the current meeting and thus promote more active discussion.*

## 1. Introduction

Meetings are a normal part of business and academia. For example, company employees regularly hold meetings for which the time and place is determined beforehand. Meetings are an important mechanism for exchanging opinions and establishing social networks. There are three types of meetings: the face-to-face type, the synchronous virtual type like a remote meeting system or chat system, and the asynchronous virtual type like a bulletin board system (BBS) or e-mail communication. We focus on face-to-face meetings because they allow written documents to be exchanged and discussed, the participant behaviors and feelings to be observed, and presentation materials to be viewed.

The advantages of face-to-face meetings justify the time and expense of setting them up. They are not limited in terms of the media used for exchanging information (written text, voice, etc.) or in terms of how they can be accessed. For instance, the participants in a face-to-face meeting about a robot and web applica-

tions under development will be more likely to express their opinions if they can touch the actual robot or operate the applications than if they can only view photographs and/or animations. If something is discussed in a face-to-face meeting, the participants are likely to understand the content more thoroughly than if they had discussed them in an asynchronous meeting.

There have been several studies on visualizing the structure of a discussion. The graphical issue-based information system (gIBIS) [1] supports the understanding, arranging, and retrieving of the remarks made during a discussion by semantically structuring them and then visualizing the structure during the course of the discussion. Matsuura [2] described a method that supports the understanding of the general stream of a discussion by structuring and visualizing the stream using surface analysis based on characters in the minutes.

We previously developed a technique, called “discussion mining” [3], in which content is semi-automatically generated from a discussion. This content can include text information, audio-visual media, and metadata. We have now developed two systems for visualizing this content. A “discussion visualizer” system visualizes the structure of the discussion in the current meeting, and a “discussion reminder” system enables the contents generated for previous meetings to be retrieved and browsed. These systems facilitate understanding of the contents of the current meeting and thus promote more active discussion.

## 2. Discussion Mining

There are two targets of discussion mining: one is to generate content from a discussion by recording and structuring the text information, audio-visual media, and metadata presented in a face-to-face meeting, and the other is to support subsequent knowledge activities through the use of the generated content. We have


been using our discussion mining technique for several years and have created several hundred files consisting of content generated from meetings.

The systems we have now developed for retrieving and browsing the contents of previous meetings and for visualizing the structure of the discussion in the current meeting support the participants in the current meeting and help to activate the discussion. This enables the participants to ascertain the divergence and convergence of the current discussion and to more equally and actively participate in the discussion.


An effective way to enable participants to ascertain the divergence and convergence of a discussion is to help them visualize the ongoing discussion. The system we have developed for visualizing the structure of the current discussion enables the participants to better understand and order the content and to contribute more effectively to the discussion. For example, a participant might transfer the right to speak to another participant who feels they have something important to contribute. If only a few remarks have been made about an important topic, the participants might try to discuss it more actively. On the other hand, if there have already been too many remarks, the participants might try to narrow the topic down.

One reason a discussion might stagnate is that the participants have unequal knowledge about the topic. Those participants who have considerable knowledge about the topic and/or are familiar with previous discussions about it can better understand the current discussion and are more likely to participate (we call them “high-level participants”). Those who do not are less likely to participate because they may lack sufficient knowledge to understand the discussion (we call them “low-level participants”). Figure 1 shows the number of remarks/statements in a year by periods that a speaker has belonged to our laboratory. If we regard a person who belongs for a short period a low-level participant, we can see that low-level participants less likely to participate. Correspondingly, these participants make fewer profitable remarks/statements. Figure 2 shows the number of agreed remarks/statements and marked ones, i.e., those that mean some participants felt were important. A way of acquiring information about agreed and marked remarks/statements is mentioned in Section 3. We can lower the threshold of discussion and thereby make it easier for low-level participants to participate by presenting information that supplements the discussion.

The two systems we present here help the participants to understand and arrange the contents of the current discussion by visualizing information about the current and previous discussions. Before describing these two systems, we briefly explain the “discussion recorder” system, which supports our discussion mining technique.



**Figure 1. Increase in number of remarks/statements by belonging period.**



**Figure 2. Number of agreed and marked remarks/statements by belonging period.**


### 3. Discussion Recorder

There has been much research into the automatic creation of meeting minutes [4, 5]. Lee [6] developed a method that records the participants’ actions and comments using cameras and microphones and then produces indexed minutes using auto-recognition technology. Chiu [7] described a system that integrates audio-visual media, text information, and presentation materials.

With our discussion mining technique, we use a “discussion recorder” system to document, retrieve, and browse remarks and statements made during a meeting. The recorder also enables each participant to reserve the right to speak. It is targeted at meetings that include a presenter, a secretary, and participants and in which the presenter presents her or his ideas using slides in a discussion room. Figure 3 illustrates an example setup for such a meeting.

The information gathered by the discussion recorder is presented on a main screen and two sub-screens. Microphones and cameras record the context information as audio-visual content.

The presenter uses a browser-based interface to display and change slides. Information related to the presentation operations, such as the begin and end times,

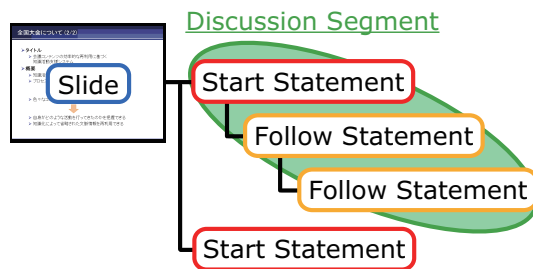


**Figure 3. Discussion room.**

is recorded automatically.

Immediately before a participant begins speaking and immediately after she or he finishes, she or he holds up a “tagging device,” which transmits his or her ID and statement type (given below). This enables the recorder to link their remarks with the corresponding scenes in the recorded audio-visual content. Participants can also use the tagging device to indicate agreement or disagreement with the statements and remarks of the presenter or other participants, to mark ones for later retrieval. The information on operation is collected and recorded by the discussion recorder. The recorder can also evaluate the remarks and identify the majority opinion.

There are two types of remarks/statements: “start” and “follow.” The start ones trigger a discussion, and the follow ones contribute to an ongoing discussion. We use these types to automatically segment the discussion to enable analysis and efficient browsing of videos. These “discussion segments” have a tree structure in which the start remark/statement is the root node and the follow ones are ancestor nodes (Figure 4).

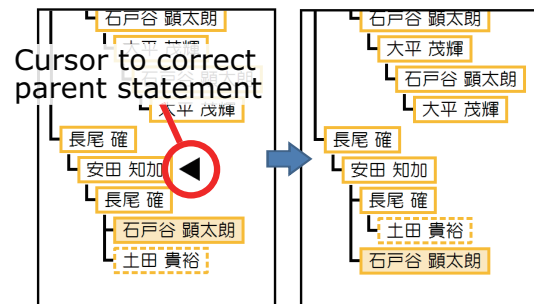


**Figure 4. Discussion segment.**

As mentioned, the discussion recorder enables participants to reserve the right to speak. A participant can clearly indicate that he or she wants to speak next. This makes it possible to clarify the order of speaking,

which helps prevent cross-talking. If a reservation is made for a follow remark when someone is speaking, the system records the reservation as a remark following the current remark and associates the reservation with the current remark. If a reservation is made for a follow remark when no one is speaking, it is associated with the last remark. The participants can also correct parent statements by using the tagging device (Figure 5).

The secretary records text information of any statements/remarks using a browser-based interface that works with the participants’ tagging devices. Any information input using the tagging devices is added to the browser-based interface as nodes. The secretary selects a node and inputs text manually.



**Figure 5. Correction of parent remark/statement associated with current remark/statement.**

## 4. Visualization of Discussion

The information recorded by the discussion recorder can be visualized in two ways: the structure of the current discussion can be visualized in real time using a “discussion visualizer” system, and the information from previous discussions can be retrieved and visualized using a “discussion reminder” system.

### 4.1. Discussion Visualizer

The discussion visualizer system visualizes the structure of the current discussion and shows it on the sub-displays. As shown in Figure 6, the discussion visualizer shows a (1) a meeting view, (2) a slide view, (3) a discussion segment view, and (4) a discussion segment list. The meeting view (at the far right) shows a preview of the participant image selected, a list of the attendees, and the elapsed time. The slide view (at the far left) shows thumbnails of the slides in order, and the one of the slide currently being displayed is highlighted. The speaker can run the slideshow by selecting the thumbnails in this view using the pointing function of a tagging device.

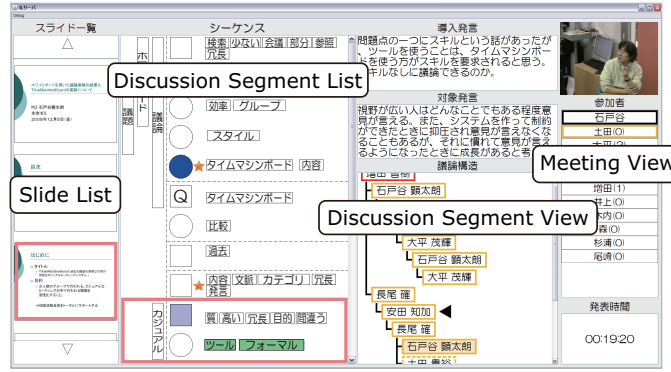


Figure 6. Discussion visualizer.

The discussion segment view (second from right) shows the text of the remark/statement that triggered the discussion and the text of the parent remark/statement of the current one (if it is a follow-type one) at the top. The discussion segment for the current remark/statement is shown beneath the text. The participants use this display area to correct the parent remark/statement, as mentioned above.

The discussion segment list (second from left) shows the stream of the discussion. The nodes representing main topics are indicated by rectangles, and those representing subtopics are indicated by circles. They are displayed as a chain structure. The keywords of multiple discussion segments are displayed on the left, and the keywords of each main topic or subtopic are displayed on the right. The nodes that often correspond to questions and answers are marked with a “Q.” The degree of agreement with each remark, as reflected by input from the tagging devices, is represented by the color intensity of the corresponding node. Icons are displayed next to the nodes containing remarks/statements marked by a tagging device. This enables the participants to identify where important remarks/statements were made.

Various types of discussion segments are created by the discussion recorder, e.g., short ones with only comments on the presentation and long ones with many remarks and statements as a result of hot debate. The long segments may have follow remarks derived from the topic of the start remark/statement, so we consider the start remark/statement to be the root node of the discussion segment and the subtopics to be derived from the root node. We conducted an observation of the subtopics acquired by the discussion recorder. We found that there are subtopics in which a participant and the presenter merely repeat questions and answers. We think that these subtopics are not important for other participants, so the discussion visualizer displays a “Q” in the nodes associated with these subtopics.

We extract a list of subtopics, defined  $\mathcal{S}$ , from a discussion segment using the procedure shown be-

low. We call these subtopics “sequences” and define  $\mathcal{B}$  as a set of a start remark/statement or follow remark/statement that has children in the discussion segment. A function called *children*( $s$ ) returns the set of remarks/statements that are children of remark/statement  $s$  (if  $s$  has no children, it returns  $\phi$ ). A function called *first\_child*( $s$ ) returns the remark/statement that is the earliest child of  $s$  (if  $s$  has no children, it returns *null*). Figure 7 shows an example of sequences in a discussion segment.

## 4.2. Discussion Reminder

A review of the contents of previous discussions can lead to shared presuppositions and common understanding of the background for the current discussion, which lowers the threshold of discussion. The low-level participants can thus participate more actively, which can lead to thinking about things from a new point of view and figuring out solutions for problems that have not been solved due to a lack of technology. Furthermore, a review will help prevent discussions redundant with previous ones. Our system for retrieving and browsing the contents of previous discussions is called “discussion reminder.”

We focused on two things in our development of discussion reminder: an accurate presentation of the discussion contents and quick retrieval of the contents so as not to hinder the ongoing discussion. Unclear and/or inadequate sharing of discussion contents can disturb flow of discussion. We therefore included a function for browsing videos of previous discussions to enable more accurate understanding.

However, review of previous discussion contents can interrupt the ongoing discussion. It is thus desirable to finish the review in almost no time by using a method for quickly finding particular audio-visual content. To enable efficient review, we added a function to the discussion reminder interface that enables the browsed information to be narrowed down such as by matching the discussion content to a query, match-

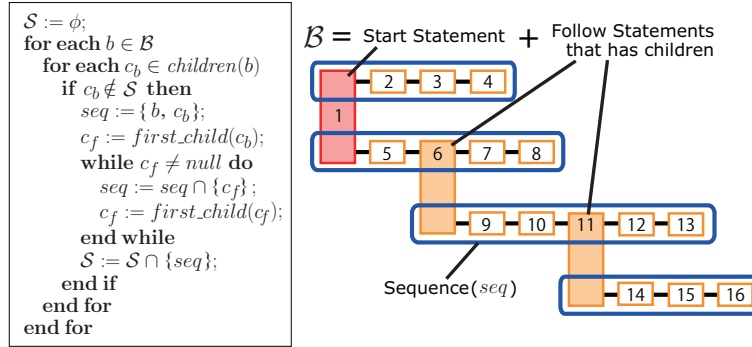


Figure 7. Sequences in discussion segment.

ing slides to the discussion content, associating statements with the matched slide, and retrieving content cooperatively with other participants.

There are two types of reminder tool, a personal one and a collaborative one. In the former case, each participant has a terminal. He or she privately retrieves a previous discussion by the terminal and indicates the found discussion to the rest of participants. The problem with using a personal reminder tool is that because a participant cannot input appropriate queries for themselves, there is a possibility of not being able to find the relevant. On the other hand, a collaborative reminder tool enables to various queries to be aggregated and previous discussions to be efficiently retrieved. We thus developed the discussion reminder as a collaborative tool.

A participant who identifies a previous discussion that he or she wants to review tells the secretary which queries to input, and the secretary inputs them on a special interface. Various types of information, such as presenter names, meeting dates, and keywords, can be used as queries. The retrieval results are displayed on the main screen, as illustrated in Figure 8. The participants can then perform various operations on these contents by using the tagging device in this interface.

This interface displays (1) a discussion content list, (2) a slide list, and (3) a discussion segment view. The discussion content list displays the titles of the discussion contents containing the discussion-matched queries. If a participant selects a title using a tagging device, thumbnails of the slides comprising the selected discussion content are shown at the bottom of the slide list. Participants can preview the larger slide thumbnail at the top of the slide list. The discussion segment view shows information about the discussion segments associated with the slide selected in the slide list, e.g., the structures of discussion segments, the speaker's ID, and remark keywords. It also enables the full text of a remark to be previewed. Participants can browse the video corresponding to the start time of the remark selected in the discussion segment view in the video view displayed on the sub-screens.

Two or more participants can search the discussion contents cooperatively. Any participant can preview the texts of remarks or thumbnails of slides while browsing using the video view. If someone selects an incorrect remark and browses the corresponding video, the other participants can indicate the correct remark by previewing the remark texts or by changing the slide selected in the slide list (Figure 9).

However, conflicts between competing operations can occur. For instance, if someone selects a remark/statement in the discussion segment view while someone browses another one using the video view, the previous browsing will be interrupted. We thus included a function in discussion reminder for excluding operations other than preview of remark/statement texts or slide thumbnails during playback using the video view until the participant who started the playback stops it or the playback ends automatically. This function enables the contents of a discussion to be efficiently reviewed without confusion.

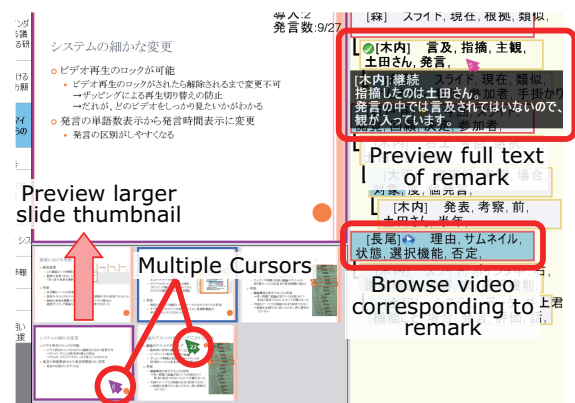


Figure 9. Collaboration in participants in discussion reminder.



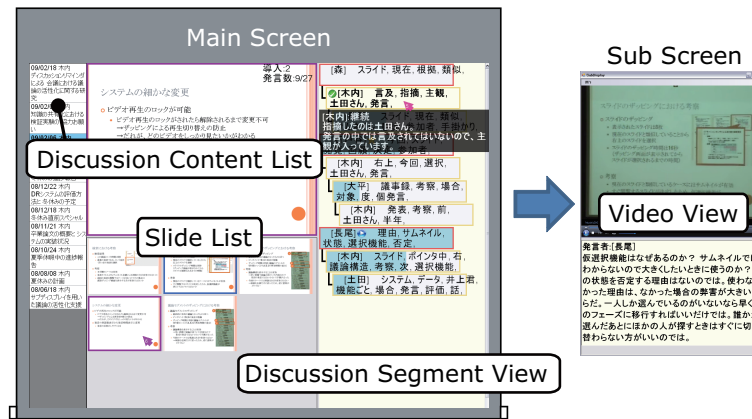


Figure 8. Discussion reminder.

## 5. Conclusion and Future Work

We have described two systems for visualizing discussion structures and contents generated by our discussion mining technique, which records and structures information exchanged in formal meetings and generates content from the discussion. The discussion visualizer system visualizes the structure of the current discussion. The discussion reminder system enables participants in a meeting to retrieve and browse the contents of previous meetings. These systems support the understanding of contents of current and previous meetings and thus promote more active discussion.

Though we did not mention it in this paper, we need to evaluate these systems. The discussion visualizer enables participants to ascertain the divergence and convergence of the current discussion by visualizing the structure of the current discussion. Thus, participants will make an efficient discussion, one in which participants break or put off a discussion about a sidestream of a presentation. The relationship between the length of discussion and presentation structure such as mainstream and sidestream will show the validity of the discussion visualizer. We will also evaluate the discussion reminder by a questionnaire survey to participants. We will ask whether they retrieved the previous discussions efficiently and whether using the discussion reminder helped them to understand the content of discussion. We think that to collect data for relevant evaluation these systems need to be operated over a long period. As a result, the number of remarks/statements by low-level participants will increase and be closer to the number of remarks/statements by high-level participants.

Our future work will concentrate on the two items listed below:

- Development of a system that enables nonparticipants to efficiently browse discussion contents
- Improvement of the systems by using information gathered by reusing discussion contents

## References

- [1] J. Conklin and M.L. Begeman: gIBIS: A Hyper-text Tool for Exploratory Policy Discussion, Proceedings of CSCW'88, pp.140-152, 1988.
- [2] N. Matsumura, Y. Kato, Y. Ohsawa, and M. Ishizuka: Visualization of Discussion Structure for Discovering and Understanding Discussion Points, Japan Society for Fuzzy Theory and Intelligent Informatics, vol.15, no.5, pp.554-564, 2003.
- [3] K. Nagao, K. Kaji, D. Yamamoto, and H. Tomobe: Discussion Mining: Annotation-Based Knowledge Discovery from Real World Activities, Proceedings of the Fifth Pacific-Rim Conference on Multimedia (PCM 2004), pp.522-531, 2004.
- [4] T. Schultz, A. Waibel, M. Bett, F. Metze, Y. Pan, K. Ries, T. Schaaf, H. Soltan, W. Martin, H. Yu, and K. Zechner: The ISL Meeting Room System, Proceedings of the International Workshop on Hands-Free Speech Communication (HSC 2001), 2001.
- [5] L. A. Rowe and V. Casalaina: Capturing Conference Presentations, IEEE MultiMedia, vol.13, no.4, pp.76-84, 2006.
- [6] D. Lee, B. Erol, J. Graham, J. Hull, and N. Murata: Portable Meeting Recorder, ACM Multimedia 2002, pp.493-502, 2002.
- [7] P. Chiu, J. Boreczky, A. Girgensohn, and D. Kimber: LiteMinutes: An Internet-Based System for Multimedia Minutes, Proceedings of the Tenth World Wide Web Conference, pp.140-149, 2001.