USABILITY OF THE MULTIMODAL MEETING BROWSER FOR REVIEWING MEETING RECORDS

BY

Karolin Blank

A thesis submitted to the School of Information Technology
in partial fulfilment of the requirements for the degree of

BACHELOR OF SCIENCE

Approved by:

Prof. Dr. A. Jameson	Prof. Dr. Dr. A. Reuter	Prof. Dr. A. Waibel

International University in Germany
Bruchsal, Germany

August 2002

Declaration

I, the undersigned, hereby declare that the work contained in this thesis is my own original work, and has not previously in its entirety or in part been submitted at any university for a degree.

Karolin Blank

Karolin Blank

Preface

During my term abroad, I studied and worked at the Interactive Systems Laboratories (ISL) at the Human-Computer Interaction Institute of Carnegie Mellon University, Pittsburgh, PA, USA. At the ISL, I was first introduced to the Multimodal Meeting Browser that has been created there for several years. Subsequent to my work, I conducted this usability study on the Multimodal Meeting Browser.

You may ask: "Well, what is this all about? I have never heard of this Multimodal Meeting Browser. So, what is it able to do and how can it help me?"

You have certainly been in a meeting before and someone usually needs to be responsible for preparing the minutes. This often is a problem, as discussions are sometimes very intense. Further problems may be a possible bias on part of the note-taker besides the unintended omission of important points. Therefore, it is difficult to have a precise record of the meeting including who contributed what ideas to the meeting.

The Multimodal Meeting Browser overcomes all these drawbacks of ordinary note-taking. The Multimodal Meeting Browser identifies speakers by analyzing their voices and their faces. Following this, the system automatically assigns the respective speaker to the speech record in the meeting transcript, as the minutes are called in the Multimodal Meeting Browser.

The Multimodal Meeting Browser encompasses even more features: Imagine you participated in a meeting and there is a follow-up on the same topic. In preparation for this follow-up, you are able to browse the relevant paragraphs in the previous meeting records and refresh your memory. Further possible activities are sending meeting transcripts off by email or listening to the recorded speech.

The Multimodal Meeting Browser has been implemented and used for demonstration purposes in the university setting for the past years. During this period, however, the Multimodal Meeting Browser was never tested rigorously in empirical usability studies.

This thesis on the "Usability of the Multimodal Meeting Browser for reviewing meeting records" includes the first usability study providing empirical data of the Multimodal Meeting Browser.

In **Chapter 1**, I will give a detailed introduction to the background, functionalities, and benefits of the Multimodal Meeting Browser. This chapter will also include the purpose as well as guiding questions and issues involved in the usability study.

In **Chapter 2**, I will first talk about the methodology used to design this study. This involves some background information, the main questions to be answered, information on the user group, the usability study procedures as well as explanations on how the data was analyzed.

In **Chapter 3**, I will present the results and evaluation of the issues that were tested. In this section, I will provide detailed information on the tasks testing an issue, on user strategies to work out the task, on user performance when they apply their strategies as well as on their comments. Following this analysis, I will conclude each of the sections and give recommendations on how to effectively facilitate usability.

In **Chapter 4**, I will give the reader a future outlook on the opportunities that the Multimodal Meeting Browser has. This information will be based on the following data drawn from the questionnaires distributed to the study participants: applicability of the system for a participant, information that users are interested in, the participants likes and dislikes with the system, suggestions for future application, the users' task performance in numbers as well as usability study rating by the users.

Chapter 5 will conclude my thesis as it will present an overall recommendation and conclusion to the usability study. This final chapter will also summarize some general observations made during the usability study.

Acknowledgements

At this point, I would like to point out the contributions and support that I had on the part of my supervisors.

I would like to thank Professor Dr. Alexander Waibel for providing me with the Multimodal Meeting Browser for my research, which was designed in the Interactive Systems Labs at the Human-Computer Interaction Institute of Carnegie Mellon University in Pittsburgh, PA, USA. I would like to thank Michael Bett and Victoria MacLaren for supporting the system's overseas installation process. I would also like to thank Professor Dr. Dr. Andreas Reuter for supporting my thesis.

Furthermore, I would like to express my special gratitude to Professor Dr. Anthony Jameson for his continuous support. I would also like to thank him for the insights he allowed me to have into the field of Human-Computer Interaction.

Finally, I would like to give my appreciation to the usability study participants for providing me with exciting material, on which this thesis is partly based:
Sebastian Hampel, Wolf-Michael Blank, Wolfgang Blank, Stefan Sulistyo, Frank Strzelczyk, Claudius Ehret, Sabine Schwöbel, Ingrid Westermann, Heike Schröder, Harald David, Radostina Georgieva, Christian Kugelmeier, Matthias Spenger, Christopher Trimble, Maximilian Mrotzek, and Olga Lamonova.

Abstract

The Multimodal Meeting Browser is an application developed at Carnegie Mellon University, Pittsburgh, USA. This complex system is able to record meetings and extract meeting minutes using speech and face recognition of the meeting participants from audio and video input. Furthermore, the Meeting Browser comprises many features for reviewing meeting transcripts and their associated audio and video recordings, such as analyzing emotion patterns, mailing the meeting transcript or searching it.

The Meeting Browser has been used for demonstration purposes, but has never been thoroughly tested with a specified user group in the setting of an empirical study. This thesis documents the first usability study on reviewing meeting records with the Multimodal Meeting Browser.

This usability study provided enormous insights into user behavior during task execution: The user works with the system using various strategies that can be influenced due to previous experience, due to internalized behavior patterns and also due to the system design.

The goal of this study is to reveal the system's achievements, its opportunities for improvement as well as its potential for future applications what reviewing meeting records is concerned.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

When thinking of a definition for a meeting, meetings are generally get togethers of two or more people that want to discuss a certain topic or issue. Furthermore, meetings usually include meeting minutes or results. These can represent, for instance, a doctor's prescription, or a final budget distribution to departments.

This plain description, however, does not take into consideration what influences a meeting and its outcome. In any meeting, multiple modalities are at play. Bett et. al. (2000) talk about "people identification", but also provide a list of "visual and verbal cues such as handwriting, facial expressions, gestures, body language, and of course speech". They believe that "recognition and integration of each of these modalities is important to create an accurate record of the meeting."

The record of the meeting is based on audio and video input streams. While a meeting is recorded, the meeting transcript is automatically extracted. This process is visible to the user in the Multimodal Meeting Browser interface. As a result, each speaker name is assigned to the respective utterance.

Identifying and assigning a users name to a certain utterance, is not easy. Bett et al. (2000) call this the "assignment problem". Certain utterances or communication cues may only be provided using a single modality. Specifically, a user may talk, but be turned away from the camera. In that case, the assignment can only be achieved by analyzing the user's voice. Opposed to this, a user may only give visual cues like nodding. Therefore, the system needs to assign this "utterance" by recognizing the face through the video input stream. In order to increased people identification robustness, Bett et al. (2000) suggest a "multimodal approach".

According to Bett et al. (2000), multimodal people identification encompasses the following components: people segmentation, color appearance ID, speaker ID, face ID,

Figure 1.1: The central navigation unit of the Multimodal Meeting Browser Interface: the navigation bar.

and multimodal information fusion.

People segmentation is used as the first step towards people identification. It is used to segment people from the background and incorporates four different stages: background substraction, noise removal, region growing and background update (Bett et al., 2000).

Following this step, color appearance identification is done by creating different models for the different meeting participants using color histograms.

Next, the speaker identification module is in charge of disclosing which meeting participant is speaking at any given time, independent of what they are saying.

Following this, Bett et al. (2000) describe the face identification procedure. This feature is only needed whenever meeting participants are dressed similarly.

Last but not least, the multimodal input needs to be combined. This information is based upon "finding the most probable configurations of people locations, identities in the room, and assignment of a speaker".

These paragraphs described how a meeting transcript is set up such that the assignment problem of who said what is solved. As described earlier, the recording of the system can be observed in the Multimodal Meeting Browser, which is the interface to the recording and browsing features of the system and which has been completely built in the programming language Java.

In my thesis, I will introduce this user interface more specifically related to the usability of browsing the Multimodal Meeting Browser. Browsing concerns benefits such as reviewing the meeting transcript, searching, editing or summarizing the transcript as well as mailing, printing or playing the whole transcript or only a selection of it.

Reviewing a meeting is greatly enhanced by the ability to employ the features named above. The meeting is bound to be more accurate than creating meeting minutes by oneself. Problems such as missing points that are only interesting for another participant, or topic bias, are omitted.

A meeting participant now is able to completely concentrate on the issues addressed and later can return to the meeting and post-process it. Then, the participant can listen to meeting paragraphs and pay attention to emotions revealed during the meeting.

Bett et al. (2000) mention the following end goals for the system that have been partly achieved thus far:

They want to implement the Multimodal Meeting Browser as a distributed application available to "participants located throughout the world". They also want to implement personalization such that the participant can "create and customize dialogue, audio and video summaries to the user's particular needs". They further would like to use the Meeting Browser for knowledge management in terms of a "database of corporate knowledge". Another goal is "quickly and accurately creating and disseminating a list of conclusions and action items". Finally, they would like to "provide rapid access to meeting records to allow browsing and reviewing existing meetings".

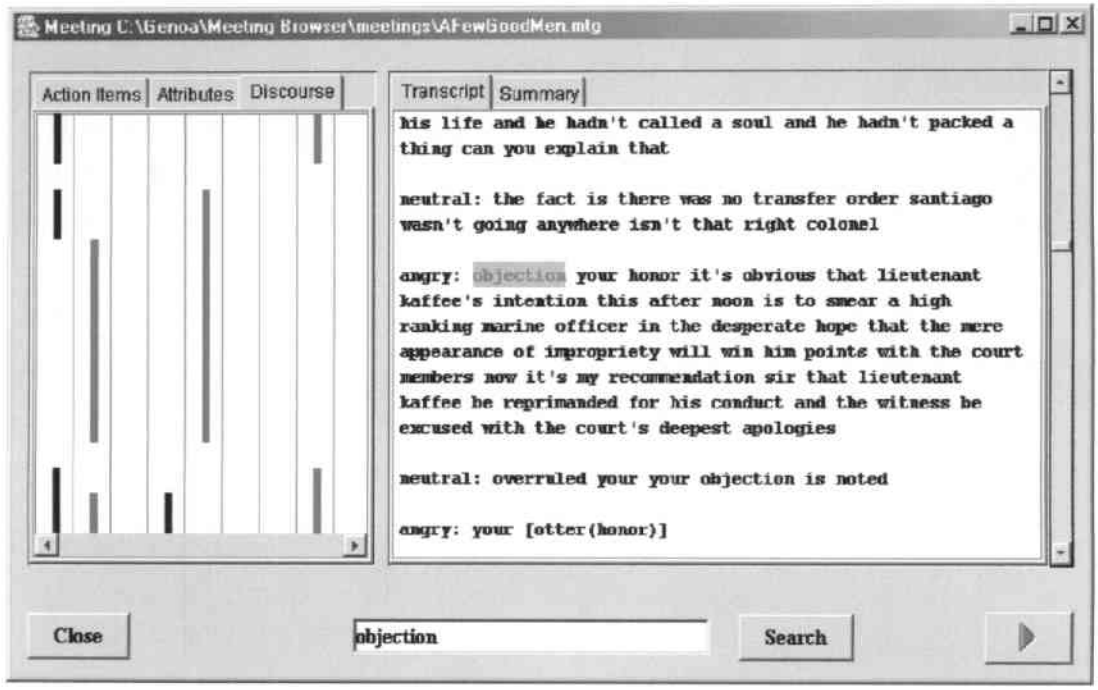

Figure 1.2: The meeting window includes the following sub windows: on the left, you see the Action Items, Attributes, and Discourse feature subwindows. The right part of this window contains the meeting transcript as well as the summaries, once they have been created. In the lower part of the screen are options to enter a keyword search as well as to play part of the transcript.

1.2 Purpose

The titel to my thesis is "Usability of the Multimodal Meeting Browser for reviewing meeting records". My thesis is therefore concerned with an empirical study on the system's browsing capabilities. This study is the first study that tests the interface for its ease of navigation or ease of learning and is aimed at revealing achievements, challenges as well as opportunities to benefit future implementations.

The design of the questionnaires, test tasks and issues to be tested for this study, was guided by four main questions (see Figure 1.1):

Table 1.1: The Four Guiding Questions

How well is the system accepted by the users?
How can the system support the user in applying various strategies for task completion?
How to teach the user to deal with system problems?
What are the opportunities for the system?

Answering these questions goes hand in hand with discussing the following eight issues (also see Figure 1.2):

Table 1.2: The Eight Issues Tested During The Study

Suitability for Searching
General Use of System
Layout and Presentation
Learning Effects on the User
Flexibility
Ease of Navigation
System Responsiveness
System Instructions

Suitability for Searching aims at disclosing how well implemented the system's search options are in the meeting explorer window, in the meeting transcript window as well as in the help feature.

The issue of *General use of the System* was generally tested in all tasks of the system and comprises small and general problems that need to be solved.

Layout and Presentation aims at disclosing problems when the user has to deal with various windows in the system, its competing layouts as well as its labels, icons and commands.

Learning effects on the user are tested once the user is facing an erroneous feature a second time. Learning effects further involve changing strategies when the user tries to achieve a certain goal.

Flexibility is concerned with supporting the user in employing various strategies opposed to forcing him into one single way to complete a task.

Ease of Navigation is closely linked to dealing with multiple window representations and its consequences.

System Responsiveness concerns issues such as time-to-load or whether a user was able to achieve a certain action using available buttons.

System Instructions were tested when it came to working with the help/online documentation feature. This issue was further examined when the user faced problems or feelings of uncertainty in other parts of the system.

Choosing these measures enabled me to create specific tasks and questions to give a thorough description of the system's lacks, and also to give specific recommendations on how the system could be improved. Furthermore, they helped me in answering the four guiding questions and gather important user information, performance, opinions and suggestions for future applications. This information is contained in the following chapters of this thesis.

Chapter 2

Methodology

2.1 Background

This thesis includes the first empirical study conducted on the Multimodal Meeting Browser. This fact puts various constraints on designing an adequate user study for this system.

Many user studies are based on hypotheses comparing two different systems or system states as suggested by Wixon and Wilson (1997). Some researchers recommend applying think-aloud methods where the user comments on the system that he is currently performing on, while others, as for instance Mayhew (1999) want their users to simply work with the system without being distracted and losing time through talking. This diversity in available methodologies does not simplify design of a study.

2.2 User Group

Defining and selecting the user group plays an important role in the test design. The user group has to be similar to the group that the system is intended for. Wixon and Wilson (1997) suggest that users must be selected along the following criteria: level of experience with computer systems, mandatory use of system, their educational background, how many other similar products they have used, number of interruptions in the environment as well as the consequence of errors.

According to these constraints, I selected 16 users from all kinds of educational and professional backgrounds. This number was a representative sample (Mayhew, 1999). These 16 users provided detailed qualitative information. This information was gathered during the test sessions, which lasted between one and one and a half hours and which included various sources for information that will be explained in the next section. Eventually, this initial study involving these 16 test users provided exhaustive data on the system, user behavior and user attitudes that I will present in later parts of this work.

I assigned all users to two major categories: 7 test users were assigned to the IT-savvy group and 9 to the IT-not savvy group. The assignment process was linked to the following criteria:

A user was assigned to the IT-savvy group if his studies or profession were strongly connected to the field of Information Technology. This group therefore consisted of one IT consultant, one IT graduate student as well as five undergraduate students with a background in Information Technologies.

Those users with little experience with or connection to IT were assigned to the IT-not savvy group. That group included three participants from the administration, one legal clerk, one graduate students as well as two undergraduate students in business. The participating civil engineer and one student of mechanical engineering were also counted for this group.

The user groups in total comprised five female and eleven male students. Three female students belonged to the IT-savvy group, while the other two could be assigned to the IT-not savvy group. The distribution for the male participants therefore was five in the IT-savvy group and six in the IT-not savvy group.

2.3 Usability Study Procedures

Each usability test was an in-field study, meaning that it was conducted in every participant's natural work environment. On the one hand, this asks for more realistic realistic results, as most users feel more comfortable in their personal environment. On the other hand, however, frequent interruptions may lead to distractions such that the test user may depend on help to succeed in the tasks.

Each user test was followed through in a pre-defined and consistent manner: The users received a warm welcome and an introduction to the Multimodal Meeting Browser. Then they were asked to complete a pre-test questionnaire gathering data on the following issues: gender, age, job title and highest degree attained, computer attitude, attitude towards learning new software, typing skill, frequency in meeting participation as well as their experience in using local and distributed applications to support meetings.

In the consecutive step, each user received some application training. Mayhew (1999) mentions that this kind of training partly represents the basis for evaluating ease-of-use of a system. In her description, a training is supposed to walk the user through the system, and provide practice to simulate expert usage.

I found further inspiration on training design from Carroll's (1990) minimalist approach to user instructions, which John Brockmann revisited in 1990. With this theorem Carroll wants to "present the smallest possible obstacle to learners efforts, to accommo-

date, even to exploit the learning strategies that cause problems for learners using systematic instructional material". Brockmann (1990) further suggests that Carroll thereby aims at avoiding learners impatience and skipping around in manuals. The benefits of the minimalist approach therefore lie in the users as they are best motivated by self-initiated exploration. I used this approach as a guideline in presenting to the participants all that they would need to know in order to get around in the system. I did not, however, give away answers. The time allotted to the application training was usually four to five minutes.

After learning about the system capabilities, the participants were asked to complete ten tasks using the system. While processing each task, the users were asked to think aloud, and they were also timed.

Table 2.1 on page 9 includes a list of the tasks given to the study participants. It further contains the purpose giving the users such a task and what issues it is supposed to test.

Those users who got stuck in one part of the question received some support in the form of hints after a certain amount of time had passed. In case the user still had problems, he received help to be able to continue. Helping users was especially important whenever consecutive tasks depended on successful completion of the previous one.

Besides designing tasks as dependent and independent tasks, I introduced two further categories: Wixon and Wilson (1997) defined result-based tasks and process-based tasks. Result-based tasks usually provide a starting point in the system as well as a goal to complete. This task setup adds some realism to the task, as the user has to structure and follow his strategy for completing a task by himself.

Wixon and Wilson (1997) proposed also an explanation for another task design: the process-based task. A process-based task outlines the steps and subtasks that are required to complete the task. Generally, this option has the advantage of collecting comparable data. One drawback, however, could be in that the task lacks realism.

Overall, there were five tasks each developed as either result-based or process-based tasks. Furthermore, six of the ten tasks depended on a previous task, the other ones were independent. Therefore, task success was an important constraint in finishing the user test.

Following the completion of the test tasks, each participant was asked to fill out a post-test questionnaire. This questionnaire sought to recap some of the system's issues, rate them and give some verbatim comments. Besides personal comments the users were asked to fill out comprehension questions, questions on their personal likes and dislikes regarding the system as well as for ideas of how else the system could be applied. Please see Appendix C on page 71 for a complete list of questions.

Table 2.1: List of Tasks 1 to 8 Including Issues Tested During Task Execution.

Task #	Task	Issues addressed
Task 1	Search on subjects specified by test conductor.	Find out about user's initial reaction to using the system and ability of system to answer such requests; Suitability for Searching, General Use of System, Flexibility, Ease of Navigation, and System Responsiveness.
Task 2	Stay with meeting and find certain info.	Ability of system to provide adequate browsing options in the system dialogue; Suitability for Searching, General Use of System, and Ease of Navigation.
Task 3	The selected text is to be extracted and sent off.	Ability of system to provide a "Mail Selection" feature; Suitability for Searching, General Use of System, and Layout and Presentation.
Task 4	Search for meeting in the meeting directory (A few Good Men).	Ability of system to support giving a clear overview on meetings available; Suitability for Searching, Learning Effects on User, and Ease of Navigation.
Task 5	Check for utterances of "objection" to an attorney's proposition in the same meeting.	Ability to make use of information received to continue with extracting further information from a specific meeting, as might occur in a browsing situation; Learning Effects on User, and Ease of Navigation.
Task 6	Connect the extracted information with available discourse definitions.	Ability of system to provide the user with discourse features; General Use of System, Layout and Presentation, Ease of Navigation, and System Responsiveness.
Task 7	Extract the system ranking and its meaning from the discourse features provided.	Ability of system to clearly explain discourse features; General Use of System, Layout and Presentation, and Learning Effects on User.
Task 8	As a check for the correct ranking play the respective utterance.	Ability of system to provide the user with play feature; Flexibility, General Use of System, and System Responsiveness.

Table 2.2: List of Tasks 9 And 10 Including Issues Tested During Task Execution.

Task #	Task	Issues addressed
Task 9	Locate the description for the "Summarize Menu" in the online documentation/help	Ability of the system to efficient online documentation/help; General Use of System, Layout and Presentation, System Responsiveness, and System Instructions
Task 10	Print the passage, in which Bin Laden's network is mentioned.	Ability of the system to give printing support; General Use of System, and Learning Effects on User.

2.4 Analysis of Usability Test Results

While the participants were completing the tasks, I kept a performance protocol and timed the users. This protocol and the timed data allowed me to come up with an objective data analysis, which I will call task performance rating.

The protocol helped me in defining the quality of the users task completion. Whiteside, Bennett, and Holtzblatt (1988) suggest that four performance levels should be set to each attribute. For analyzing the participants' task achievement, I therefore created four levels of task quality:

Using the optimal method or one deviation gave the user 100% in the quality level. Two deviations still were worth 75%, while four to five deviations or hints left the participant with 50%. In case the user had more deviations or got completely stuck such that he depended on help to complete the task, he received 25% (see Table 2.3).

Table 2.3: Task Achievement Rating

100%	Up to 1 deviation
75%	Up to 3 deviations
50%	Up to 5 deviations or hints
25%	More deviations or help

Task performance does not only depend on task achievement, but also on time. Carroll (1997) already described that users become impatient when they are not quickly rewarded when performing a task. This fact largely influences the evaluation of the task performance rating. It may be that a person performing a task with little or no deviation also performs it in a very short time.

I argue that it is also possible that the user spends a lot of time before completing a task, or that a user spends a lot of time but is unable to complete a task without help by the evaluator. These factors need to be taken into consideration, and therefore both values need to be combined.

The calculation is achieved by standardizing the fraction of task achievement and the time needed to fulfill the task, see equation 2.1.

$$\textit{Task Performance} = 1 - (1/\frac{\textit{task achievement}}{\textit{task time}}) \tag{2.1}$$

Besides the objective measures, I analyzed subjective data. This data comprised participant comments during the task execution, subjective rating of features and issues applicable for the Multimodal Meeting Browser, as well as subjective data from the comment section in the post-test questionnaire.

Chapter 3

Results and Evaluation of Issues tested

In this chapter, I will present the problems and findings collected in the usability study. These I will organize by the issue tested with the Multimodal Meeting Browser. Each section will be further divided into tasks testing the issue, user strategies and results, subject comments on this feature as well as a variety of recommendations and possible solutions to cope with these problems.

3.1 Suitability for Searching

During the usability study, the participant faced two search options in the Multimodal Meeting Browser: The search feature in the meeting explorer window as well as the keyword search in the transcript window containing the meeting minutes.

3.1.1 Tasks Testing This Issue

In Task 1, the users first experienced the explorer search option in order to located and load a meeting. In order to narrow down the meetings available, they knew that one specific user participated in that meeting while another one did not. Also in Task 4, the users had to use the search feature to continue. This depended, however, on how the users found the results in the previous search.

In Tasks 2, 3 and 4 the users further had to apply keyword searches in the meeting transcript to be able to continue completing the respective tasks.

Tasks 1, 2 and 3 have been dependent on each other. In order to reach the second search feature in the meeting transcript window, it was important to succeed in the previous task and find and load the correct meeting.

3.1.2 User Strategies And Results

Opening the Meeting Explorer

Participants generally approached the task using one of two methods. In Task 1, 3 of the 16 users chose the menu bar for opening the meeting explorer. The remaining 13 participants chose the open icon. This number decreased when the users had to open the explorer a second time: merely participant *F* chose the menu path for this subtask. This may have been due to the fact that I used the icon during the training session and only pointed towards options available in the menu bar.

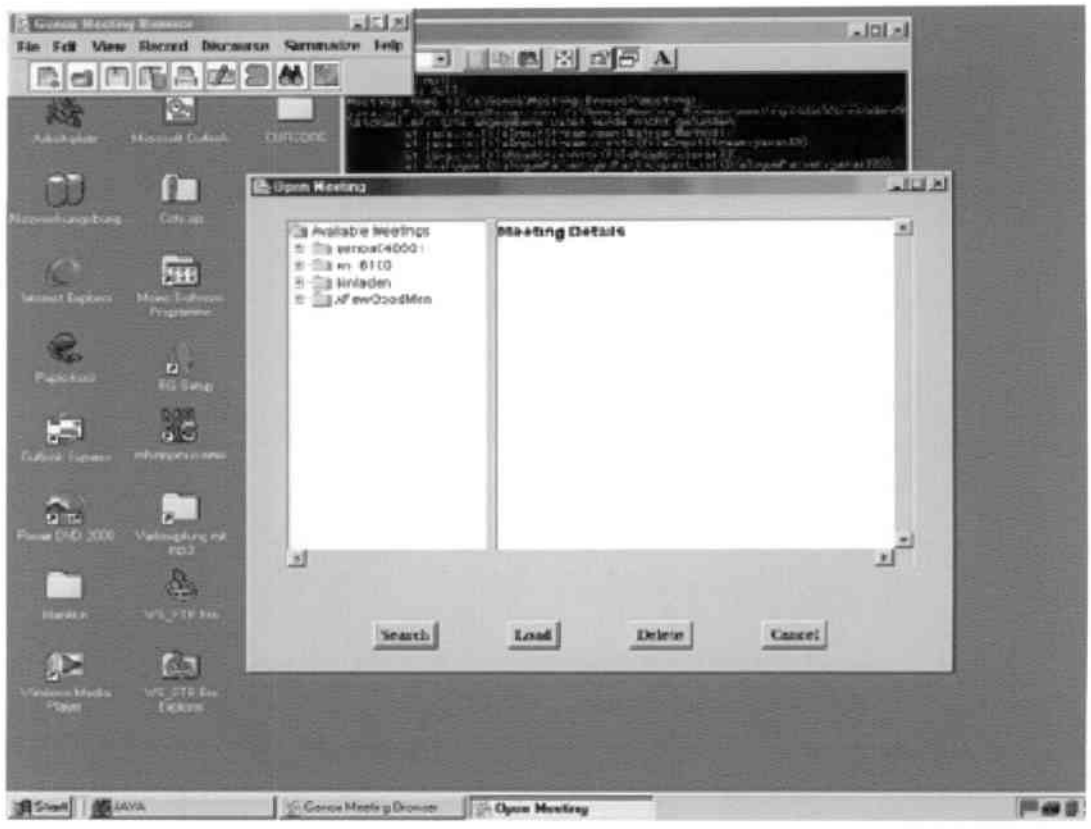

Figure 3.1: When the user loads the Multimodal Meeting Browser and opens the explorer window, the screen will be covered with three windows for navigation: the navigation bar, the explorer window, and the dos shell displaying the actions happening in the background.

Finding the Correct Meeting

The subtask of finding the correct meeting sounds straightforward, but there exist several paths for succeeding in that task. The optimal strategy is clicking the search button in the lower part of the meeting explorer. This opens the search window (see Figure 3.2 on page 15) where the user were supposed to enter "A" with connector "but not" to the right and enter "Chad" in the consecutive line. The meeting "`en_6100`" would have been the only match for this query.

By *examining the properties available* some users attempted finding the correct meeting. Even though this was not part of the training the participants received, they seemed to be very attracted by the properties associated with each available meeting.

In Task 1 the participants were asked to find a meeting where a person called "A" participated but "Chad" did not. Users *C*, *D*, *J* and *L*, succeeded finding the correct result with the help of the participant properties. But they were not too sure about their findings as the participant properties show different representations, i.e. "A, B", "Two Participants".

Therefore, many tried to located and use the search option. Users *E* and *L* had actually problems finding the search button in the meeting explorer. It seemed to be out of their view when they kept looking for it in the menus in navigation bar and the icons (see Figure 1.1 on page 2). The other participants, however, were more goal-oriented and followed the strategy presented during the training:

By clicking the search button in the meeting explorer the users opened the search window and perform the search. At this point, many users encountered problems assigning logical connectors to the participant text fields, knowing that there is one participant called "A", but that there is no a participant called "Chad". I will now explain how the users approached the task opposed to the optimal strategy:

A number of 6 participants applied a search on "a" or "A" only. Another user assigned the correct "A" and assigned "but not" to the right of Chad in the consecutive line. Both strategies returned two available meetings to the meeting explorer. All 7 users then succeeded in excluding the wrong result. This procedure was respectively chosen by 3 IT-savvy and 4 IT-not savvy participants.

Many users are accustomed to using search engines in the internet, which are based on topic searches. 5 participants, 2 IT-savvy and 3 IT-not savvy, therefore stuck to applying a search on a topic. They did not receive any results and lead them to the fundamental problem of the "Blank Search" that I will explain later.

There were two more situations when the system did not return any results. This can attributed to their prior knowledge dominating performance explained by Bhavani and John (2000): One further group of users did not picture the assignment as a chain of

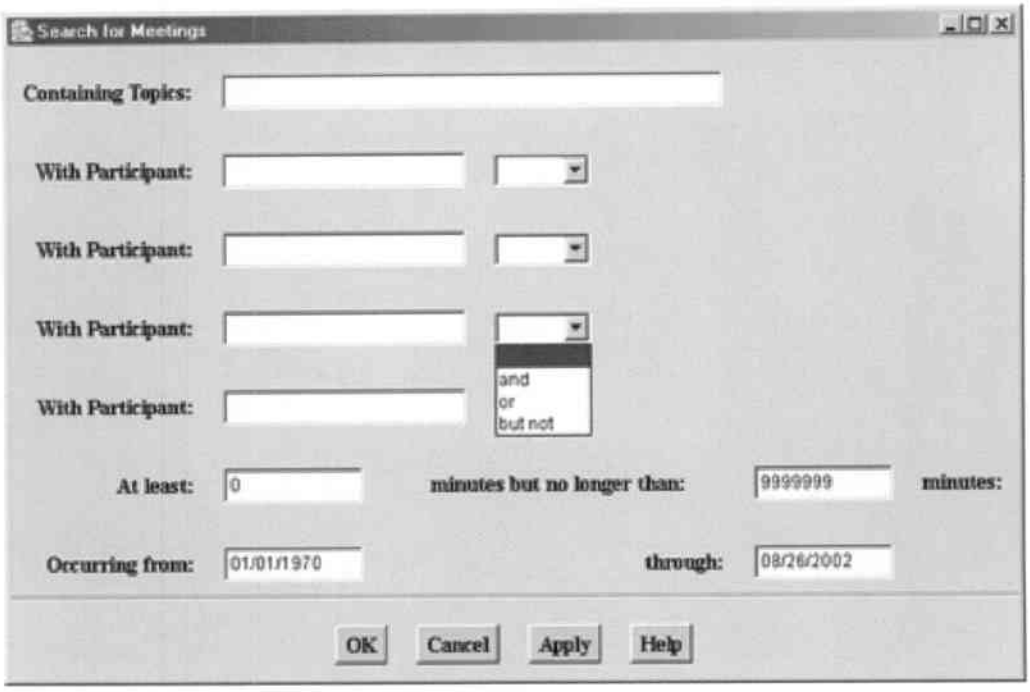

Figure 3.2: This window illustrates the current implementation of the search window.

participants and logical connectors. They rather assigned the connectors from right to the left, and thus regarded each line as a separate assignment. It is interesting to see that these users, users *B*, *C*, *E*, and *P*, were experienced in using computer systems, and that three of them had used the SAP R/3 system before. They were therefore accustomed to applying the logic required for that system when they set up the query. Another group, users *F* and *P*, dealt with a similar problem, they only switched the lines. Both also had an IT background.

Nonetheless, one third of the participants or 5 users, users *A*, *D*, *G*, *I* and *L*, applied the optimal method for retrieving the correct meeting. User *D* was one of them, and was the only one who did not take up this result, even though he had earlier found the correct answer already in the properties.

I will now explain the problem of the *"Blank Search"*. All participants that did not succeed in retrieving the correct result or two result options were required to apply a "Blank Search". That means that the user has to apply another query from the search window. This search has to leave all fields empty that are available for entry. Applying only this search will return the complete list of meetings available in the meeting explorer. As this kind of search does not work in an intuitive manner, the participants were taught

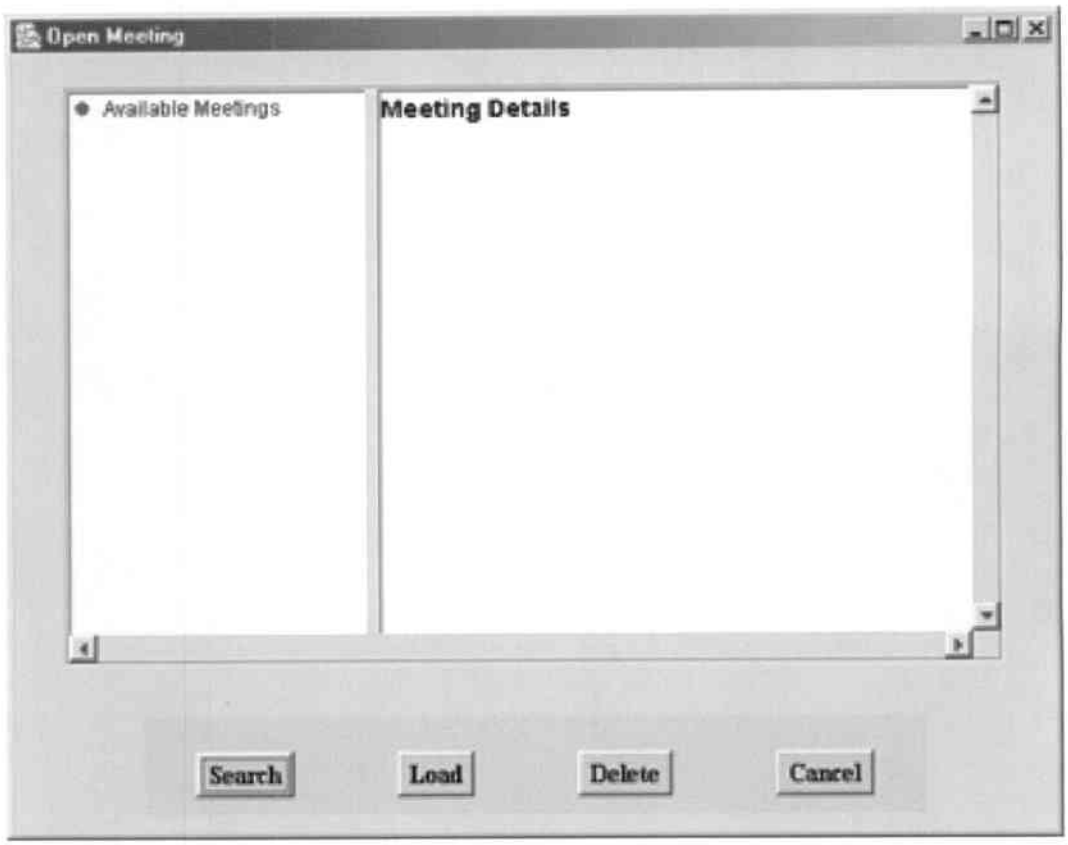

Figure 3.3: The empty explorer window is the only information for the user that no results were returned. How to continue from here, has not been supported yet.

how to deal with the "Blank Search" at the first encounter. From the second encounter on, dealing with this issue was taken into consideration for the user's performance and learning advancement.

Searching the Meeting Transcript

In Tasks 2, 3 and 5, the users were asked to find certain keywords in the meeting transcript as part of the question.

The keyword search proved to be full of obstacles for most users. 8 users, mostly IT-not savvy ones, attempted to scroll and read the meeting transcript in order to find the result. Only when realizing that this is not a very efficient method they took a closer look at other options available in the window. Only then, they used the keyword search on the bottom of the window.

Having the correct intention did not always help users to achieve their goal: user *H*

made a typo while user *C* hit caps on the keyboard and entered "BUNGEE" instead of "bungee". Both queries did not return any results in the text and left the participants discouraged.

It usually suffices for many search engines to hit "Return" on the key board to initiate a query, and 5 users applied this procedure repetitively with the transcript search. They did not succeed. Clicking the search button therefore was the only way to initiate their query.

If the user wishes to find another result, the user can click the button another time. This option was intuitively enjoyed by many users. It was nonetheless bound to produce problems: User *H* double-clicked that button and unconsciously omitted one result important for completing the task.

Another problem, I observed, was that many users seemed to reach the end of all possible search results: Having found the last result, the search simply stopped and there was no way to return to the first result and start over or go on from there. Therefore, most of them tried to find the initial result by scrolling upwards through the transcript. As this was not a very efficient method, some users closed and reopened the meeting window to start over again. The IT-savvy people seemed less bound to encounter this problem: 2 IT-savvy users, users *C* and *G*, opposed to 5 IT-not savvy ones, participants *J*, *L*, *N*, *O* and *P*. Users *N* and *O* even dealt with this problems three times, which was also reflected in their task performance.

3.1.3 Subject Comments

Concept of "But Not" in the Search Window

This concept was mainly critiqued by users that were used to the R/3 system. User *C* mentioned that SAP has introduced a true/false system which allows separate assignments to single objects and rows to the left. Therefore, he was inexperienced in using logical connectors in the way it was implemented here. Still, he was able to work with them.

Blank Search

The participants were mainly concerned with giving suggestions to avoid this situation: User *A* would have liked to see a "Show all" button for the meeting explorer and user *E* suggested adding a "Clear results" button to the search window.

Keyword Search in the Meeting Transcript

User *G* would like to see a "Next" button to make searching for the next result more obvious and available. She would also like to see search results including content similar

to the query, and user *N* would like to extract indirect requirements. User *F* explains this in more detail: He mentions that with an XML set up of this window, it would have been easy to find connected nodes on activities. In his point of view, this would have greatly enhanced search speed and finding results. *F*, however, did not give any detailed description of how this could be implemented.

User *K* has a completely different approach to this problems: She tries to extract keywords by looking at the color lines. In her point of view the color lines represented shortcuts to topics and she did not see any immediate connection between the transcript paragraphs and the color lines.

Participant *C* mentions further problems: he played a paragraph and then wanted to search the transcript for a keyword information. This however, was no longer possible. He had to close the meeting and reload it in order to be able to continue. User *E* encountered the same problem in the opposite order: *E* searched for an item first and then played the paragraph from the transcript. When he then wanted to continue searching for another item, he had to also close the window and start over again.

Continuing After Reaching the End of Search Results

At first, user *J* was lost and not sure how to continue. In his point of view, having an information box would be a good solution. He suggested a text message such as "Finished searching, would you like to start again?".

Results from the Post-Test Questionnaire

On a scale from 1 (worst) to 7 (best), all participants were asked to rate the system for its ability to search meetings. The average of the IT-savvies was at five and the average of the IT-not savvies was at 6, giving an overall average of 5, 5.

3.1.4 Recommendations

In the following, I would like to recap the problems users had when they searched the Multimodal Meeting Browser. I will also provide recommendations on how to solve these problems (see Table 3.1 on page 19):

Blank Search

Eliminating this flaw in the system will enhance usability of the search feature in the explorer window. Furthermore, implementing buttons would assure the user in what options are available: a "Show all" button in the meeting explorer window, and a "Clear Results" button in the search window.

Table 3.1: Recommendation on Suitability for Searching

1. Eliminate "Blank Search".
2. Give feedback to user actions and provide solutions for problems encountered.
3. Support indirect search on topics related to the query. This could be done using overview+detail procedures such as color highlighting and change in font size as in the Popout Prism application.

Continuing After Reaching the End of Search Results

The users must be able to continuously browse through their results. Similar solutions were found in other text processing systems like Textpad and the Microsoft applications. This could be achieved by adding "Next Result" button or arrow symbols to the meeting window.

Furthermore, it must be ensured that users are able to issue search queries and play meeting parts in an alternating manner.

Keyword Search in the Meeting Transcript

This concerns the topic of "how-am-I-strategically-looking-for-results"? Many users proposed having a keyword search that was also capable of searching for similar topics. For such a feature, certain color representations dissimilar to the ones used for the discourse window could implemented. Highlighting the keyword result green in the transcript, for instance. Depending on the closeness to the actual keyword, similar topics could then be highlighted in colors with growing difference to green if their predicted closeness is growing bigger.

There is an existing system called Popout Prism developed at Park, which focuses on overview+detail applications and analysis in the field of Human-Computer Interaction. In their paper, Suh, Woodruff, Rosenholtz, and Glass (2002) describe methods of highlighting, changing font size and applying different levels of transparency. Furthermore, they use the Microsoft Explorer for this purpose and created two subwindows: on the left, there is the overview part, which contains the complete document as well as a box signifying the document excerpt displayed in detail in the right part of the screen.

This concept could be carried over to the Multimodal Meeting Browser transcript window, simply by adding another "tab" to the subwindows on the left in order to display the overview. For the purpose of showing keyword results, the detailed result information would then be integrated in the actual transcript. By doing this one can provide detailed

search but still maintaining an overview on the document. interface.

3.2 General Use of the System

This section is concerned with how the users dealt with the system in general. I will demonstrate how participants were able to cope with features previously unknown to them.

3.2.1 Tasks Testing This Issue

The general use of the system was tested actually in every single task. There are some tasks that present the best insights into general problems associated with the Multimodal Meeting Browser.

Tasks 1 and 4 were representative tasks for loading a meeting. Tasks 2 and 3 delivered important insights on the general use of the transcript. Tasks 6 and 7 displayed reactions to the color lines implemented for the purpose of mood and sentence tracking. The final task that gave me insights into the general use of the system was Task 9. In this task, I watched users work with the play feature of the Multimodal Meeting Browser.

3.2.2 User Strategies And Results

Loading a Meeting

In the user training, participants observed the evaluator in loading a meeting twice. The optimal method was demonstrated, which is in opening the meeting explorer, highlighting a meeting and clicking the load button in the lower part of the window. In the observations, many users were severely distracted and only few participants applied the optimal method right from the beginning.

The users were asked to load a meeting for the first time once they had found the meeting in task one. Some users were stuck in the properties (see Figure 3.4 on page 21) or looked into the navigation bar for the load option. Of 7 users, 2 IT-savvies, *B* and *C*, opposed to 5 IT-not savvies, *G*, *H*, *K*, *M* and *P*, double-clicked the meeting name for loading. Some others got stuck in the properties available when the users tried to open the meeting by double-clicking items in the properties.

Most of them were distracted for a moment but then realized by themselves that there had been an alternative way for loading a meeting. At this point, they often lost some time by thinking about how to continue, but were generally able to use the highlight-click method.

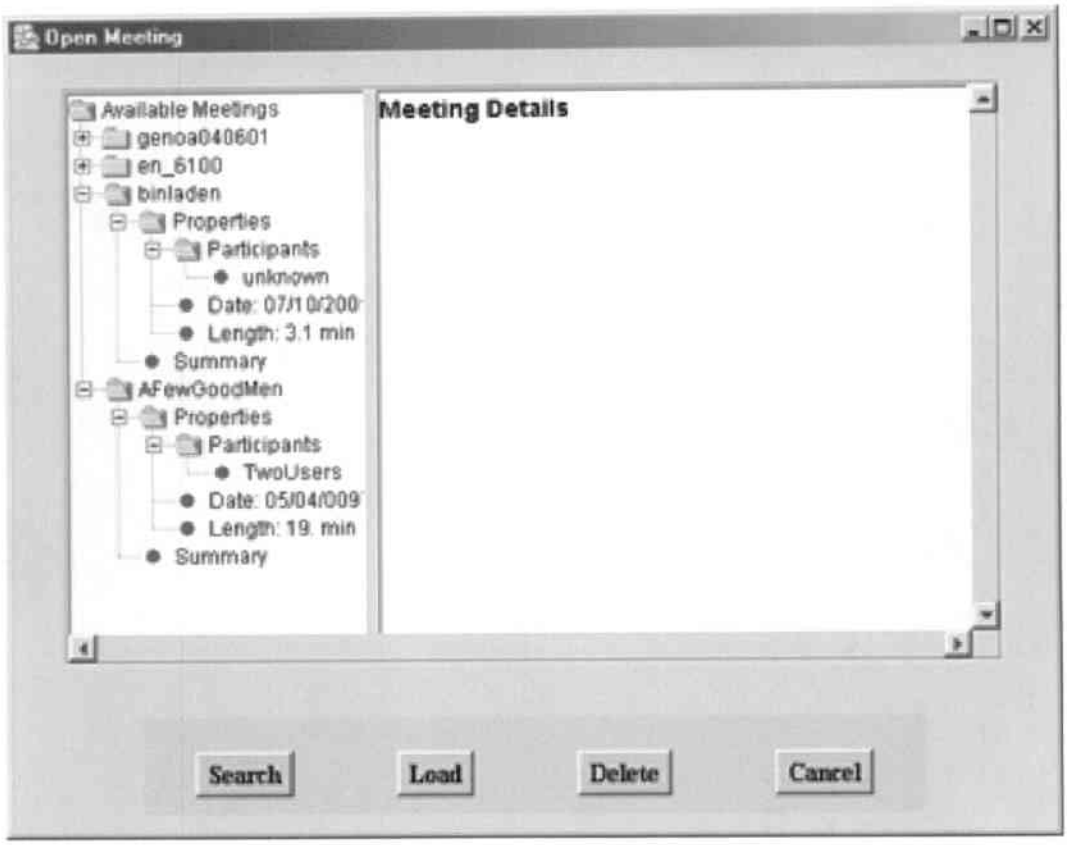

Figure 3.4: This view on the explorer window gives insight into the properties setup and labelling.

General Use of the Transcripts

Once the user succeeded in loading the meeting, he was again confronted with a new layout and a new window. There were several issues that influenced user performance:

Users *B* and *E* tried to get a better view on the meeting transcript by maximizing the meeting transcript window. Even though this action enlarged the window as such, the objects contained did not change their sizes. Only minimizing helped the user to refocus.

Moving around the meeting window was not always easy. Some users were forced to use the scroll bar: even though user *P* attempted to scroll by placing the cursor in the transcript and using the down key on the keyboard, she had to arrange herself with using the scroll bar.

Color Lines in the Discourse Window

In Tasks 6 and 7, the users were asked to investigate the color lines in the discourse window and define meaning to them. Again in these tasks, the participants applied different approaches:

User *A* believed that each color line represented the speech track for one single user when he looked at them.

Users *I*, *K*, *L*, *M* and *O*, among them only 1 IT-savvy participant, defined meaning to the color lines by sliding over them and taking a look at the info box.

Some other users, users *A*, *B*, *G*, *H* and *N*, tried to define meaning to the color lines using the text paragraphs assigned to them. Again, this strategy was taken up by only 1 IT-savvy and 4 IT-not savvy participants.

Clicking the lines seemed to be the obvious strategy for users *E*, *F*, *H*, *J* and *M*.[1] With the help of this action, they discovered the properties window. This properties window contains more precise information on the transcript paragraph for the users.

The IT-savvies were also well represented in the group of users that played the paragraphs associated with the color lines. Users *C*, *D*, *K* and *P* chose this method to find more information about the color lines and the text.

The Play Feature

There are two possible ways that were considered as optimal for playing a certain paragraph: first, the user highlights the desired length of the transcript and hits the "Play" button in the lower right corner of the meeting window. The second option is playing the paragraph from the properties window.

There was no training on this task. It is, however, clear to see that 14 users started playing a paragraph from the properties window. This may be due to the fact that they were asked to investigate the properties window in an earlier task.

It was interesting to see that one user, user *K*, only found the play option when she tried to close the properties window. Others saw that feature in an earlier task and were very eager to playback the track. Only two users, users *B* and *E*, succeeded in playing in the meeting transcript.

User *C* had been successful in playing a paragraph at an earlier instant using the properties window. Later, though, he attempted to replay the paragraph from the transcript by placing the cursor and hitting the play button. As he did not make a selection, however, he was unable to replay the paragraph using this strategy.

[1] Some users were not sure about their definition found using a previous strategy. Therefore, they are listed in another section as well.

Even though the users enjoyed playing the paragraph over and over, they did not realize their strategies' constraints: Using the properties window restricts them to playing only one sentence or paragraph. Highlighting any length of the transcript and hitting the play button, gives the user much more flexibility.

After playing a paragraph for many times, users *F*, *J* and *L* attempted to stop the playback. This, however, was not directly possible. Using the "File/Close" option on the player was the only means for stopping the file.

3.2.3 Subject Comments

Loading a Meeting

User *L* was somewhat irritated by the naming convention used for loading a meeting. "Open Meeting" is used for both, the icon and the menu item in the navigation bar. Clicking either option first opens the meeting explorer. When actually opening a particular meeting transcript, the users are asked to load. This wording is not very straightforward for most users. User *L* therefore would rather like to see an "Open" button instead of a "Load" button in the meeting explorer.

As mentioned earlier, many users were irritated by the properties displayed in the meeting explorer and also tried to find loading options there. User *J* therefore asked for adding a category for direct access to transcript called "Text": "I would expect a function "Text" among those available when I click on a meeting. Currently there only is "Properties" and "Summary". Loading the meetings text via the "Load" button instead of double-click is not a normal procedure".

User *A* further talks about the inefficiency of the properties conventions: "naming, and listing are unclear". When taking a closer look at the conventions used for listing meeting participants, the user is confronted with a list of names, but also with the simple statement of "Two Users" or "unknown". This is not a very effective way for the user to determine whether he would like to enter this meeting without using the search option.

General Use of the Transcripts

The transcript formatting also led to some confusion. One meeting included transcript notes such as `#human-noise#` whenever a speaker breathed or another indistinguishable noise was recorded. Users *C*, *J* and *O* paid attention to this problem. While user *J* was amused about this utterance, user *O* mentioned that this feature is not intuitive to understand.

Some participants also commented on transcript formatting. Users *K* and *P* felt restricted in understanding and applying their strategies when it came to understanding the

meeting's contents: there are no capitalized letters, wrong spelling, no sentences and the meeting names are neither easy to remember nor easy to distinguish.

Color Lines in the Discourse Window

The participants' comments on the color lines were very diverse: some users were almost euphoric about them. Others did not even want to take a further look at them, and it was very hard to get them to do the tasks on these lines at all.

User *H* was the most prominent example for this: *H* absolutely disliked the color lines. She believed that these lines would not mirror the content of the paragraph and that she felt irritated by the lines. Therefore, she encountered severe problems in dealing with the task, and was almost stubborn about it.

User *I* defined meaning to the color lines on her own. Therefore, she came up with her own definition. Red, in her opinion, was something positive, and blue had another color coding also. The setup of the colors to her was not logical, but confusing. User *J* tried to define meaning to the colors according to a certain pattern, but did not succeed. User *K* had a completely different idea of the lines, as she believed that the lines represented levels of intensity.

User *A* argued that the lines were not self-explanatory, especially since the same color was used twice. User *D* almost would not have considered the blue line for evaluation, as it begins in a different paragraph. Also to user *C*, the line's size was unclear when he compared it to the the paragraph's length.

For user *O*, it is not clear what the discourse window reflects: "Does it only represent one paragraph, or does it refer to everything displayed in the right window?"

Many have further suggestions on how to improve the feature: User *A* mentions that "colors and columns were not explained in any part. So, maybe beams of speakers could additionally be displayed in the columns. This could facilitate a visual search for the contributions of a certain meeting member."

User *O* said that the discourse feature window had to be more accurate and explaining. User *P* also believes that the discourse features should include some differentiation on who says what.

On the Properties Window

User *H* believes that the lines would be unnecessary for the ordinary user who is only interested in the transcript and not in emotions. What concerns the properties window, she mentions that such details are of no value to the end user. The only thing of interest would be the participants as well as start and end times. She also mentions that the intensity scale is not obvious to her. After finishing the task, she tries to close the window

using "x", which does not work. Almost every other user in the group encountered this problem, though. Among them user *J* who further disclosed that the close option "ok" only works when the top bar of the window is highlighted. Instead of being able to close the properties window, opening it can also be achieved by right-clicking the color lines.

User *L* is not sure about the meaning of the properties window's contents. User *J* also is unsure about the meaning of properties. He defined the properties as "background info of each line". He tracks back this definition to a friend's psychology thesis.

Last but not least, user *H* would like to see a scale for measuring the level of intensity. Right now, she is not even sure about what the intensity measures.

The Play Feature

After critiquing the color lines and associated features, user *H* is delighted about the play option. She admits that there is some advantage to the lines, though only when listening to an audio file: "Very Hollywood!"

Results From the Post-Test Questionnaire

General use of the system by labels and commands was rated 5 on average with no significant difference between the IT-savvies compared to the IT-not savvies.

3.2.4 Recommendations

The following section will explain solutions and further extensions in more detail. Please see Figure 3.2 on page 26 for more detail.

Loading a Meeting

First of all, a consistent properties layout is needed, which is filled with consistent information.

Second, the "Load" button should be adapted to an "Open" button.

Finally, it could be useful to integrate another load option by double-clicking within the properties listed in the explorer window.

General Use of the Transcripts

It would be good to implement some typesetting procedures to post-process speech after recording. This would greatly enhance readability and understanding of a meeting's topic and content.

Table 3.2: Recommendation on General Use of System

1. Implement a consistent properties layout for the explorer window.
2. Include the option to load a meeting by double-clicking from the properties.
3. Use typesetting for the transcript.
4. Differentiate and explain colors and color meanings.
5. Selecting a piece of text and the playing it needs to be more obvious. This could be achieved by a good help feature or manual.
6. The discourse properties window was of nearly no use to the participants. It therefore needs more useful information for the user browsing a transcript or needs to be eliminated.
7. The categories listed in the discourse properties window need to be explained.
8. Change the explorer window's "Load" button to an "Open" button.

Furthermore, navigating the transcript should be enhanced by adding the option of scrolling through the transcript using the cursor.

Color Lines in the Discourse Window

It should be made more obvious what each line track represents. Adding headers would also prevent users from confusing the index tabs with headers for the discourse feature tracks.

Color repetition does not enhance usability of the discourse window. Highlighting the text assigned to a specific category in that color besides adding more colors would be one solution to this problem.The users did not find the category for sentence or YN-quest[2] very obvious. This was also due to the fact that some syntax categories were falsely assigned or not at all mentioned. Another problem was that some lines could not be accessed.

Some users had problems regarding overlapping color lines as two separate ones, and did not consider them for evaluation. Clearer differentiations between single paragraphs are therefore necessary to support user understanding. However, this may be hard to implement as two voices are sometimes interwoven, which then stretches the color line over several paragraphs.

[2]YN-Quest is the abbreviation for Yes- or No-question.

The Play Feature

It was not clear that parallel lines contained the same audio file. The users expected some category specific utterance instead.

Except for the play feature, the properties window did not seem to be of much use to the users.

Another important aspect currently hindering user actions is to get the properties window close button to work.

The users need more advice on how to use the play feature in a more efficient way. As we have seen in the above examples, selecting text, though easy, was not an obvious and familiar strategy. Bhavani and John (2000) suggest that efficient strategies must be made explicit and that there may be a weak causal relation between method and quality of the product.

3.3 Layout and Presentation

This section will describe how the users performed using the different window representations. Layout and presentation are important factors in learning how to use the system. Users are looking for general patterns in layout and presentation, which helps them to navigate more easily through the system.

3.3.1 Tasks Testing This Issue

In fact, most tasks dealt with the problem of coping with various windows and their respective layouts. Tasks 3, 6, 7 and 9, however, are best to demonstrate issues related to layout and presentation.

Task 3 will be exemplary for how the user mails a selection to a friend. Tasks 6 and 7 will again show another aspect of the color line specification. Task 9 will then refer to the layout and presentation of the help feature.

3.3.2 User Strategies And Results

Mailing a Selection

Task 3 required the users to select a certain paragraph and then send this selection by email. The optimal method for this task includes selecting a paragraph following by entering the file menu in the navigation bar and then clicking "Mail Selection".

Selecting the correct part of the transcript was not difficult for any user to achieve. It was not obvious to them, however, that the system remembers the selection made last,

even if the selection is visually lost.

Next, there was the subtask of mailing this selection to a recipient. Participants using a menu-based approach for tackling the problem were in a great advantage. Users *I* and *J*, were not distracted by any other buttons or options. They used the correct approach and succeeded in fulfilling the task using the "File/Mail Selection" procedure very quickly.

Figure 3.5: The "Mail Selection" option is hidden in the navigation bar.

All other users, however, encountered severe difficulties, on the one hand, due to prior knowledge dominating performance (according to Bhavani & John, 2000), and due to the "Mail Transcript" icon in the navigation bar, on the other.

In total 8 users, 3 IT-savvies, users *C*, *E* and *K*, and 5 IT-not savvies, users *J*, *L*, *M*, *N* and *O*, attempted to right-click and save the selection. They then wanted to mail the saved selection to the recipient using other email software like Microsoft Outlook. Users *C* and *K* were persistent in that logic as they also checked out the Edit menu for achieving their purpose. It was not obvious to them that such a software is integrated in the Meeting Browser itself.

The users' confusion can best be illustrated by the various attempts they made to find

the "Mail Selection" option. Users *A*, *C*, *K*, *O* and *P* consulted the Discourse Feature window as well as the Attributes and Action Items windows for that feature. Users *O* and *P* clicked Summarize Meeting. User *D* played the paragraph again.

Finally, some users stuck to the navigation bar and the icons in particular. The "Mail Transcript" icon seemed to be the most applicable to the users since all users except 3, users *C*, *I* and *J*, clicked this icon. As this icon calls the complete transcript in the mail window, this obviously can not be the correct option for completing the task.

Letting the users proceed and discover the problem by themselves, led to the following discoveries: Of those 13 participants only 4, users *A*, *G*, *K* and *L*, realized by themselves that this option cannot be correct. User *J* excluded this option by merely looking at the description. It is interesting to note in this context that 4 participants of those discarding the "Mail Transcript" option were not-IT savvies.

Not succeeding by using the "Mail Transcript" icon, however, led to some more confusion on part of the users. They very much stuck to the icon options and did not consider the menu items in the navigation bar for a long time. Some were quite frustrated and therefore needed hints on where else to look.

Once having drawn their attention to the menu items in the navigation bar, many still had to look for a long time to find the "Mail Selection" item, which is actually located in the lower menu list of the File menu. This proved to not be a very intuitive location for any mail option. The users usually started out looking into the Edit menu and then moved to the right in the available menu options. In case they accidentally entered the File menu, the users generally did not move or look down far enough to see any Mail option. Once found, however, they oftentimes chose "Mail Transcript" as this option was listed first and "Mail" seemed to be the only keyword they were looking for.

In the end, all participants were able to finish the task. Comparing task execution times as well as the number of deviations to the other tasks, however, the average task performance was the worst for this task.

The Multitude of Windows

The current setup of the Multimodal Meeting Browser involves over seven different windows, each having a different layout. Using a multitude of windows that sometimes include even a variety of different subwindows, makes using the Meeting Browser very confusing and oftentimes distracting. There are: the navigation bar, the meeting explorer, the meeting window including five sub-windows, the properties window, the mailing window, the help/documentation window, the play feature window and more. In this section, I would like to introduce various constraints that apply when working with these windows (see Figure 3.6 on page 30).

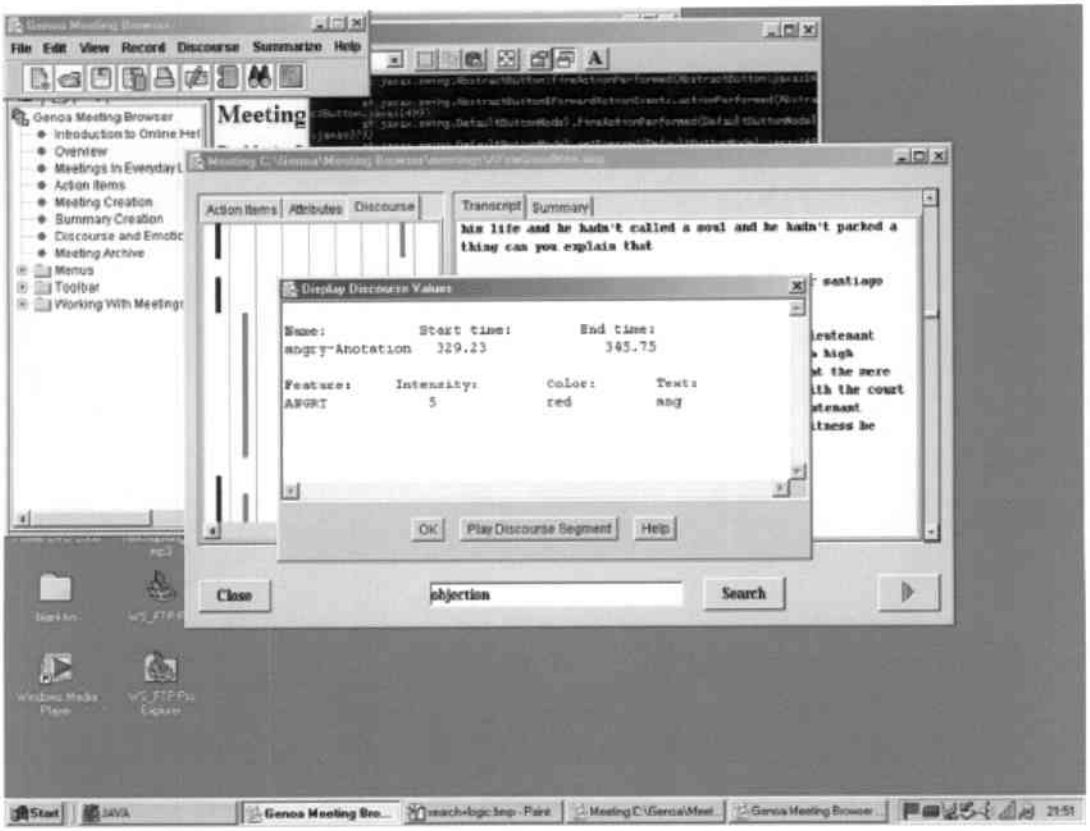

Figure 3.6: Numerous windows do not support easy orientation, the properties window on top.

Participants kept looking in the wrong window for information that they expected there. User *C* was only one of those mistaking the Action Items sub-window for a header in the discourse feature window. This lack of clarity distracted participants sometimes so much that they forgot what they were asked to do in the beginning. One prominent example occurred when the users were looking for the Mail Selection that I have explained in more detail in the previous section.

I can further say that even though the navigation bar is meant to be the center of navigation in the Meeting Browser, this does not apply for all features. Some users work menu-based and expect to be able to begin an action starting from the navigation bar. Some features, however, are not supported by the navigation bar. This applies for the play feature, for instance.

There were some inherent problems involved when referring to the properties window. This window cannot be closed by clicking the close button in the window frame. Closing can therefore only be achieved by clicking the OK button. It is important to also note

in this context that this option for closing is only possible when the window frame is highlighted blue.

It has some advantage that the meeting explorer window is supposed to be closed after loading the meeting itself. This however, often is not the case and the window remains open for no specific reason. When the window closes itself, however, the user has to find some new orientation as he does no longer have a direct visual connection to neither the location nor the path that he used to load the meeting. The user is virtually lost in space.

Let me now talk about the scenario that the user would like to compare two voice streams using the properties window of a certain color line. This intention leads to six open windows on the screen: the navigation bar, the meeting window, two properties windows and two players. Generally speaking, the words are visually highlighted once spoken. In case the user now attempts to watch the text while he plays a paragraph, the windows are blocking the view (see Figure 3.7).

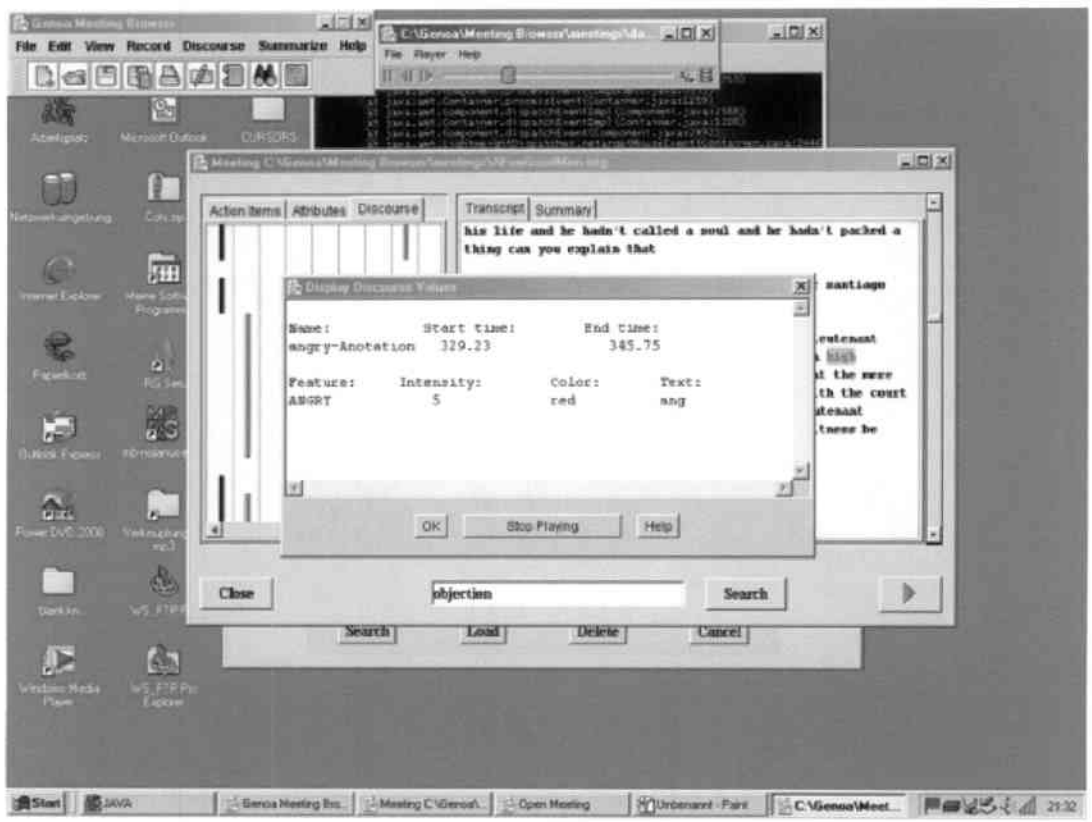

Figure 3.7: Playing a paragraph using the properties window, leads to a playback where the view on the transcript is blocked.

Another problem can occur when the participant closes the meeting, but the properties windows and the respective players remain open. There is no more connection to the meeting transcript available any longer. Therefore, all open windows connected with the previously examined meeting transcript should be closed as well.

Due to this problem, some users attempted closing those windows by themselves. Unfortunately, they closed one window too many when they closed the Meeting Browser itself. This happened to users *C*, *J* and *M*.

Help/Online Documentation

Task 9 examined using the help feature of the Multimodal Meeting Browser. Here, users were asked to extract information on the "Summarize Menu". The best strategy to find such information is by choosing the help menu and then "Online Documentation" or by clicking the icon with the question mark in the navigation bar. In the following, the best strategy included entering the menu folder in the left part of the screen and retrieving information on the Summarize Menu. The method was perfect if the user compared what he read to the actual menu available in the navigation bar. This approach to understanding the content was only chosen by users *K* and *M*.

This time only 3 users, IT-not savvies *A*, *I* and *O*, relied on the icon, while the rest called the help feature using the menu. 6 users including 5 IT-savvies, users *C*, *D*, *E*, *G* and *P*, and 1 IT-not savvy, user *A*, then chose to read the first article sounding similar to what was asked for. They were content reading information on the text "Summary Creation". Only user *G* realized that this information does not give her any news on the Summarize Menu.

The layout seemed to be quite standard in the first place, but the users had to cope with some errors: The text formatting of the Summary Creation article did not wrap to match the window size. Therefore, those users had to maximize the screen to view the whole text. Another problem that user *L* mentioned was that he was unable to position a cursor in a text field to use the up or down buttons for scrolling purposes.

As a last aspect on this issue, the search feature integrated in the help section is important. Users *B*, *K*, *L*, *N* and *P* were 5users that would have liked to rely on the search option provided but not functioning in the Help/Online Documentation interface. Improvements in this part are certainly appreciated from both sides, IT-savvies and not-IT savvies.

3.3.3 Subject Comments

Mailing a Selection

User *K* and *M* mentioned that the arrow integrated in the mail window looks like a send button. As it did not initiate the expected action both users were irritated.

User *A* did not see any menu header related to "Mail" and therefore had problems finding the "Mail Selection" option in the navigation bar. Furthermore, he mentioned that "marking text and sending mail in two different windows, which are far apart, does not make task achievement easy".

User *C* believed that this is a very nice feature to creating and updating meeting excerpts. He, however, confused the summary icon with "Mail Selection" mentioning, that "I was confused about finding the "Mail Transcript" icon, but [I did not have any] connection to the other mail options. My suggestion in this instance would be adding another mail menu tree to the menu items. This could also be integrated into the icon itself by prompting the user for whether he wants to mail the selection or the complete transcript when clicking the icon."

Pasting any selection, if the user has actually made one, should in user *D*'s point of view be the default setting.

User *E* would appreciate the implementation of the following: "I would like to have [the] right mouse button [available] within the transcript window, which includes features such as email, print, copy and so on." User *C* referred to his experience with the SAP R/3 system mentioning that he can automatically mail a paragraph by right-clicking a selection and then choosing a mail selection option.

In general, I may say though, that the users were quite impressed by this functionality after mastering it. User *O* also mentioned this fact that it "takes sometimes quite some time", but he rates the functionality in general "useful".

Participant *M* considered the system suitable for mailing selections of the transcript "due to the fact that the utterances are clearly distinguishable from the different meeting participants."

The Multitude of Windows

In the post-test questionnaire, the participants were asked to comment on how they felt about dealing with the different windows involved in the Multimodal Meeting Browser.

User *C* had a window problem and argued that "orientation is not facilitated due to the different layouts of each window". Nonetheless, he believed that he would get used to working with them. User *B* shares this attitude, but has further critiques, such as: "The windows seemed to be too small. The windows background or any other application

running in the background is distracting. Unneeded windows should be hidden, i.e. the Dos-shell at system startup."

Participant *L* was one of those feeling lost when the meeting explorer closed itself, once he loaded a meeting. He was not sure whether he had chosen the correct meeting and therefore wanted to check back with the available meetings in the explorer. He mentioned that a new version should "support navigation between features. Otherwise, the navigation is basically a matter of luck for you to get where you want to go." User *L* was not alone with his problem, as user *P* also would have liked to see a button allowing to go back to the initial explorer window.

User *D* mentioned that "the properties windows were blocking his view when watching the the words being played from that window such that he first had to move the windows before being able to extract text info required in that task".

User *F* felt that the feature information was hard to find in the unorganized properties window.

User *O* mentioned that the discourse feature window has to be more accurate and explaining. And participant *P* critiqued that there was not made any differentiation between the speakers in the discourse window.

Participant *C* further made a remark on the discourse window layout: "it is unclear to me how to differentiate between the traces in this window as I confuse the labels for different windows with each trace." Furthermore, he had his own idea of how to solve the problem of the different windows involved at this point in time: "I believe that all windows should be integrated into a big whole. Right now there is no integration of all features/windows and no common menu structure. Therefore, it is necessary to find a common solution. What concerns the pop-up windows, these should also be integrated into the big picture and should hold the same menu structure as the main navigation bar."

User *E* agrees with the one window approach. For him, one advantage of such a solution would be being able to see the icons at all times. He also discovered the problem that maximizing the meeting window does not look very nice. Also, participants *G*,*L* and *M* would prefer a more integrated look of the Meeting Browser windows in order to be able to keep track of the contents.

There were a couple of users, users *I*, *J*, *O* and *P* that were content about how the system is currently organized. User *J* only encountered some problems in using the help feature. And user *P* felt the need of more explicit graph representations. She would like to see the development of emotions throughout the meeting as well as the emotional development respective to the user.

Participant *F* was confused by the multitude of windows as well as the fact that he had to close one to get to another. User *K*, however, realized that using these windows helps

users in multitasking, taking part in or browsing different meetings at the same time.

Help/Online Documentation

While proceeding with the task, a couple of users took the chance to talk about their impression of the help system layout:

Figure 3.8: The same bullets are used in the explorer window (left) and the help feature (right). The only difference between both is that the bullets in the explorer window do not contain explicit information, which leads to misunderstandings when working with the help feature.

Participants *A* and *C* mentioned that they would not have expected any content in the listed items in the left part of the screen as the dot was also used in the meeting explorer, however, not containing any information (see Figure 3.8 on page 35). *A* came up with this reason as he did not see the cursor change to a hand. When he checked out the window, he had to scroll a lot. Therefore, he would like the help window to open at a bigger size on the screen.

Referring to the info received by reading the Summary Creation article, user *A* critiques that it does not explain much. The sample summary given also does not suggest that it actually is a summary. Nor does it tell him how big the meeting used to be prior to the summary.

The participants also seized the opportunity to make remarks during the post-test questionnaire.

User *E* said that he would like to see the help function work in order to be able to better cope with errors.

Participant *J* also discussed the format that he expects a summary to be like: "The sample summary in the online help does not meet my understanding of a summary." A summary should be neutralized, i.e. along the lines of "A stated that..." "B pointed out...". He believes that direct speech is not efficient since this makes it very difficult to recognize it as a summary. In addition, "the summary was almost impossible to read since absolutely no punctuation was used."

Results from the Post-Test Questionnaire

On a scale from 1 (worst) to 7 (best), all participants were asked to rate the system for its quality in layout and presentation as well as its ability to mail a selection.

Regarding layout and presentation the average of the IT-savvies was at $4,3$, including one rating of 1, and the average of the IT-not savvies was at 5, giving an overall average of $4,625$.

Regarding the Meeting Browser's ability to mail selections, the average of the IT-savvies and the average of the IT-not savvies nearly even, giving an overall average of $6,094$. This feature received the absolute top rating in comparison to all other features. Interesting to note since the user performance rating was very low.

3.3.4 Recommendations

The recommendations will be explained according to the layout categories they refer to. Figure 3.3 on page 37 contains an overview of such recommendations.

Mailing a Selection

Bhavani and John (2000) mention in their second visualization strategy that it is necessary to view relevant information instead of irrelevant information. Therefore, it must be made transparent to the users what options are available for a certain piece of content.

First of all, it is necessary to add another menu list to the navigation bar clearly marking the options available for the mail categories. These choices should only be accessible when the user has actually opened a meeting. This should also be implemented for the summarize option, as well as all other options that should not be available until a meeting is loaded.

Table 3.3: Recommendation on Layout And Presentation

1. Add a menu for "Mail" options to the navigation bar.
2. Allow right-click options while dealing with the meeting transcript.
3. Limit user confusion by integrating the multitude of windows into one solution.
4. Allow to change the meeting window's size for getting a better view on the transcript.
5. When closing the meeting window all other windows associated with the transcript should also close.
6. In the help feature, focus on what the user needs to perform his work and get the search options to work.

Moreover, adding further navigation options by allowing right-click selections would provide the experienced user as well as the system with valuable shortcuts. The content of the right-click menu should also be adapted according to the user action. For instance, if the user makes a selection in the meeting transcript the available options could be "Play Selection", "Mail Selection", "Show Properties" and "Copy".

When nothing is selected and the user opens the right-click menu, however, the available options should include the following: "Play Meeting", "Play this paragraph", "Mail Transcript" etc..

The Mail icon can further be connected to a system prompt asking the user what he would like to mail, the complete transcript or simply a selection. This feature could also be made independent from prompting by defaulting to Mail Selection when the user has highlighted some text in the transcript.

The Multitude of Windows

In my point of view, the best approach to solve the problem of multiple windows would be by incorporating the windows into one big window solution.

This approach has several advantages and challenges:

Having an overview over what options are available at all times as well as quick navigation between features, are only two advantages. Such an organization of the system features could also be used to facilitate better system visualization. This can be supported by another strategy suggested by Bhavani and John (2000). Navigation in global view supports manipulation in local view. This concept is, on the one hand, based on the prop-

agation of important system information that includes making dependencies known to the computer as well as exploiting dependencies to generate variations. On the other hand, this concept depends on the system organization, which involves making dependencies known to the computer and generate new representations from existing ones.

Integrating the windows in one solution, however, probes the designer with new challenges. This specifically applies to integrating the various windows. One challenge is retaining the option to have several meetings open at the same instant. The opportunity in this case is that new ways of representation can be used and applied in order to make it even more obvious what the properties represent and what message they are attempting to convey.

Help/Online Documentation

It would be helpful to users to find information on the how-to for functionalities listed before the user receives background information. The help feature is an essential part in communicating knowledge to the user. It is supposed to convey hints on how to deal with problems, but not give away all the answers.

Bhavani and John (2000) featured information on identifying strategic knowledge, which has to be achieved before one can tackle the problem of promoting efficient aggregation strategies in a system. Besides training and manuals, help systems and tutorials play an important role in teaching the user to work efficiently. This is also one possible way to avoid user frustration.

Therefore, in the next version of the Multimodal Meeting Browser the help system needs a lot of improvements compared to the current status. One of the necessary changes is adopting the task of redesigning the content list. Using the same list icon in the help feature as in the meeting explorer did not prove to be an effective measure. The users did not expect relevant contents opening such links. Therefore, choosing a different icon may be necessary.

Further issues that must be considered in the future concern the intensity rating of the speech properties, the "Blank Search" if not eliminated as well as bringing the keyword search to work.

Seizing the chance to create a good help feature will definitely motivate users as they will be able to perform tasks quicker.

3.4 Learning Effects on the User/Result-based tasks

This section will be concerned with explaining user learning while using the system. Learning can best be mirrored giving users result-based tasks, which have the following characteristics:

According to Wixon and Wilson (1997), a result-based task provides the participant with a starting point to a problem as well as a description on what to find. Such a task setup requires the user to come up with his own approach to solve the task.

Learning success and task success are closely linked to focussing and motivating users. John Carroll et al. (1997) mention five trade-offs to summarize a psychological design rationale view of minimalism. Among others, they mention that errors can be frustrating and disrupt task goals. Therefore, diagnosing and recovering from errors help the user to focus and self-motivation as well as to sharpen a concept of correct performance.

3.4.1 Tasks Testing This Issue

Applicable for this sections are Tasks 4, 5, 7 and 10. They are all related to loading a meeting from the meeting explorer, using the search field in the meeting window and/or involve the application of a "Blank Search" (see Figure 3.3 on page 16).

3.4.2 Adapting User Strategies And Results

Loading a Meeting

Loading a meeting seems to be an easy task to complete. Highlighting and loading, that seems to be all. However, in the first task the users faced even more constraints on this issue.

It was quite interesting to observe, how the general user strategy changed with growing experience in using the system. At first, most users were acting in an exploratory way and often got lost on the way to task completion. Later, having acquired more experience, their strategy changed to a more goal-oriented approach.

A very prominent mistake, in the first attempts, was double-clicking the meeting name in the explorer window to load it. Other problems included being stuck in the properties of the meeting, looking at the navigation bar for a loading option or simply loading without highlighting the desired meeting.

Later, most users omitted any redundant action in the meeting explorer. Only 3 users compared to 7 in the first attempt encountered deviations or problems when loading a

meeting. Among those 3, 2 did not encounter any difficulties loading the meeting in the first attempt.

Comparing the user behavior of the first two attempts to the third one, 6 participants again tried to load the meeting by double-clicking it. Only 2 of them, however, made this mistake for the second time.

Regarding the task of loading a meeting, I see improvements in user strategy. Even though all users were taught in using the correct method, many participants, however, were unconsciously using another concept and double-clicked the meeting name instead of highlighting and loading it.

Applying Searches

I described user behavior related to searching in an earlier section. Now, I would like to explain how users improved in adapting their strategy during the study.

The "Blank Search" was a big problem that best should be eliminated in future versions of the Multimodal Meeting Browser. At first encounter, the participants were taught how to deal with this flaw, and even though many were not happy to work with it, they learnt how to deal with it in future task scenarios.

Depending on the search the participants entered in Task 1, they had to face the "Blank Search" either right away or when they had to load another meeting. 8 users dealt with the "Blank Search" in Task 1, 5 of which had to deal with the "Blank Search" again in Task 4. Only 2 users then needed help recovering from this status. It is true that the concept of the "Blank Search" is quite hard to grasp, for the second attempt many users first had to think about how to do it. After a while, though, they were able to come up with the correct strategy on how to return to the complete list of meetings.

Nonetheless, implementing a better solution than the "Blank Search" is completely necessary. This can be achieved by adding buttons called "Show all" in the meeting explorer or "Clear" in the search frame.

What concerns searching the transcript text, however, users were unconsciously hitting the return key rather than clicking the "Search" button next to the search field. It seems like the IT-savvies were more inclined to hitting the return key than the IT-not savvies. Nonetheless, after the third test, 2 IT-savvies and 3 IT-not savvies used the search button. However, 3 of these users correctly used the search button in the earlier tasks.

I therefore argue that there is no direct link between learning and an effective strategy in this matter. The participants rather unconsciously decided, which option to choose. This could also have been affected by diminishing concentration as the participants performed this third in-text search in the final task.

Printing a Selection

Another one of Carroll et al.'s (1997) five trade-offs was proven in the context of printing a selection. They defined this trade-off as follows: "Working on familiar task orients and motivates learners by engaging prior knowledge, but may encourage task-specific learning and engage inappropriate prior knowledge."

Printing a selection is a task known to us from tasks like printing a selection of text from a web page. You make a selection, use File/Print and click the selection option in this print window. Nothing else had to be employed also during this task scenario.

During the user study, participants encountered the task of mailing a selection. Therefore, in this context, users seemed to know how to proceed with this task. This time, however, the users had to find another strategy to achieve a solution. Many users were struck with uncertainty when they were unable to find any category called "Print Selection" - neither in the meeting window or the menus nor in the icons. This happened to users *A*, *C*, *E*, *L*, *M*, *N* and *P*.

Due to their uncertainty, users were prevented from exploring the tool further. Many considered clicking the print button, but discarded that idea before taking it into consideration again. Finally, 12 users, 5 IT-savvies and 7 IT-not savvies, used the File/Print option in the navigation bar, while 4 users, 2 users each, decided to click the print icon from the navigation bar.

3.4.3 Subject Comments

Applying Searches

User *A* mentioned that he needed assistance from the supervisor when it came to applying the "Blank Search". Otherwise he would not have found back to the meeting overview.

User *K* would have liked to see more obvious hints for the user that applying a "Blank Search" is now required.

What concerns searching the transcript, user *B* realized that hitting return was no option for initiating a query.

Furthermore, participant *N* found that he was only able to search for one keyword at a time in the transcript window. Therefore, he had to adapt his search strategy by issuing more simple queries instead of one including many keywords.

Printing a Selection

Only user *M* mentioned that he had a problem using the print option. In fact, he needed a lot of guidance during the task execution as he was focused too much on the "Mail Selection" idea.

Results From the Post-Test Questionnaire

There was no specific issue to rate regarding the user's learning evolution during the usability study. Therefore, the information had to be drawn from the study protocols.

3.4.4 Learning Synopsis

In essence, this section suggests that learning has been successfully inferred and strategies adapted by the users for the tasks of loading a meeting and searching for meetings. What concerns searching within meetings, the users did not apply a rigorous strategy and rather worked by unconsciously making decisions.

Finally, printing a selection taught the users that newly learnt strategies are not always transferrable to other task scenarios. Adapting their strategy was the only consequence for their actions, and a familiar procedure, i.e. printing a selection from a web page, was discarded.

3.5 Flexibility

Flexibility plays an important role in strategies. As described in earlier sections, users tend to have their own strategies to achieve a certain goal. The system must support those users and allow them to choose many paths to succeed in their purpose.

3.5.1 Tasks Testing This Issue

Flexibility was tested during Tasks 1, 8 and 9. Task 1 tested the options available for users to load and search a meeting. This has been discussed in detail in earlier sections. In this section, I will therefore explain flexibility using the example of playing meeting excerpts, Task 8.

3.5.2 User strategies and results

Playing a Meeting Excerpt

It is obvious that the participants have some degree of flexibility using the play feature. There are two possible ways of running the play feature: Firstly, participants can select as much text as they wish and then hit the play button at the lower corner of the screen. Or, secondly, they can click a color bar related to a paragraph in the transcript and start playing there.

The discrepancy that users generally started the play feature from the properties window instead of selecting and playing the text, depended on two things: first, it was not

evident to users how to proceed in order to play the feature from the transcript window. Second, the users had to use the properties window for a previous task and therefore saw or even tried out the play feature in that window.

There is one point that would speak against the second option. Some users unconsciously tried to apply both ways, as happened to user *C*. He chose to play a paragraph using the play option in the properties window when he was not explicitly asked to do so. Later, he seemed to have forgotten about that option when he was explicitly asked to play. At that point, he tried to position the cursor in the transcript and hit the return button. The system, however, did not give him sufficient flexibility to chose this route. As this experience diminished his motivation, he did no longer explore the possibilities and try to play after selecting a piece of text, but needed hints to go on.

User *E* applied the same method, however continued to explore and finally run the play feature from the transcript.

Other Applications of Flexibility

The Multimodal Meeting Browser provides more examples giving the user sometimes more, sometimes less flexibility.

The "Blank Search", for instance, is one feature that I have given a lot of attention to already. It is therefore clear to see that the "Blank Search" does not give the user lots of freedom in applying different strategies. The user is rather forced into one single procedure.

Flexibility also covers the subject of dealing with the transcript. It is true that there are many features available for use in the navigation bar. However, when the user is currently working with the meeting transcript, the navigation bar does not represent a direct access to the functionalities. Making access available through the right-click would definitely enhance flexibility and user satisfaction.

3.5.3 Subject Comments

In the post-test questionnaire, all participants were asked whether the system gave them sufficient flexibility in working the way they wanted to. The following is a collection of their comments:

User *D* mentioned that "most functions could be reached through different channels, however, it still forced a certain way of thinking and working on the user. Once accustomed to this way, it was easy to extrapolate the rest, though". Participant *M* agreed to some extent, as he felt like having to deviate from his strategy.

User *G* talked about adding a keyword search that would query on topics related to the object of interest. Such a feature could help users learn about related topics.

Participant *K*, however, felt very appreciative of the systems capabilities: "it was pretty flexible concerning all the options, like printing or mailing specific paragraphs, and the possibilities to track previous meetings".

Results From the Post-Test Questionnaire

Besides the comment section on flexibility, all participants were also asked to rate the level of flexibility that the system gave them on a scale from 1 (worst) to 7 (best).

The average of the IT-savvies was at $4,3$, including one rating of 2, and the average of the IT-not savvies was at $4,7$, giving flexibility an overall rating of $4,5$.

3.5.4 Flexibility Synopsis

There was only one user, participant *A* who was not very fond of the idea of flexibility: "I don't need any flexibility, but one clear way of how to achieve my tasks." In his point of view, too much flexibility would cause confusion. I can only agree in part to such a conclusion.

While it is true that many experienced users apply their knowledge for processing a task in a schematic manner, it is also true that many users have many different approaches to completing a task.

There is a famous slogan in election campaigns: "You need to catch the people where they are." This idea should, in my opinion, be adapted to the use of any software. Serving different user strategies means a lot more work regarding implementation, however, it will add to customer satisfaction, usability, user acceptance and eventually the success in the market.

In our case, this would aim at eliminating the "Blank Search", bringing the right-click option to life, and allowing users to employ their own strategies instead of forcing one single solution on to the user.

3.6 Ease of Navigation/Process-based tasks

Ease of navigation is yet another issue that was tested during the usability study. Ease of navigation can best be tested by giving the users process-based tasks:

Wixon and Wilson (1997) have defined process-based tasks as those where "steps and subtasks are outlined for the user by the evaluator". One disadvantage to such an approach is that the tasks do not represent real situations as they are rather constructed. Nonetheless, the advantage is that the collection of comparable data is possible.

Ease of navigation means that the user can work in a flexible manner and apply his strategy for completing a task. This is another crucial aspect of task success and again raises the question of how the user can best be supported to achieve this.

3.6.1 Tasks Testing This Issue

Tasks applicable to this issue are Task 1, 2, 4, 5 and 6. Task one will help disclose how navigation affected task performance. Task 2 has the clear goal that the users are to extract transcript content. In Task 4 then, navigation was tested by observing user behavior in the light of switching back and forth between the explorer window and the search mask. Analyzing Task 6 allowed to become aware of the fact that navigating through the color lines was not always easy.

3.6.2 User Strategies And Results

Switching between windows was one important factor for evaluating navigation. When the users worked with the system for the first time, they had to get used to switching between windows. Problems in this regard were not related to moving between windows. Problems were rather connected to concentrating on the one window that they were dealing with. In case other windows had to be taken into account, they did not pay attention to them but rather stuck to the first window and tried to solve the task using that window alone.

As the system was new to all test participants, many were exploring during task execution. This attitude sometimes distracted them from performing the tasks. Loading a meeting is only one example: The users diverged from the optimal and taught method once they realized that there was more information hidden in the explorer window. This fact led some users to being stuck in the meeting explorer and to losing time.

Focusing within the meeting window and paying attention to the transcript only is also difficult to achieve. The users need some more specific eye-catcher to have a starting point. This could be done by setting the transcript font color to 100 percent saturation, but

setting saturation of the accompanying windows to 50 percent or 60 percent saturation.

In Task 6, while defining meaning to the color lines in the discourse window, all IT-savvies and one third of IT-not savvies were inclined to clicking the color lines for further information. They explored and navigated without being yet asked to do so.

There was one special case of participant *H* who felt unable to adapt to the system and who was extremely confused when switching between or within windows. This may have been due to the fact that the in-field testing was in her work environment as the assistant of administration. During the study, she received six phone calls and seven people came to see her. She was practically multi-tasking and switching attention from visitors, over calls, and to the Multimodal Meeting Browser. She then forgot what she had learnt in an earlier task. It was quite difficult for her to refocus when paying attention to the Meeting Browser. She clicked from one window to another, reread the task and needed a while to recap what she had done and was just about to do before she was interrupted.

3.6.3 Subject Comments

In the post-test questionnaire, all participants were asked two questions that are directly linked to the issue of navigation:

First, the users were asked whether the labels, options and commands on using the system were easy to follow and use. And, second, they were questioned whether it was easy to move around the different parts of the system.

Switching between windows hindered user *B* in moving easily around the system in the beginning. He called this confusing.

Users *C* and *D* felt that they would adapt to using the system with a short time of practice. User *D* further mentioned that he liked the common menu frame that was visible all the time. Both users therefore awarded a 6 and a 7 to that issue. Participant *L* felt like he would get used to dealing with many windows, even though he would prefer one single window.

User *F* had a different attitude towards the issue of navigation. He said that moving around the different parts of the system was "not always easy because of too many open windows". He therefore rated this feature with a 3.

Participant *G* had problems finding some functions in the first place. She reasoned that these problems could be attributed to the unprecise description of the buttons in the explorer window. Later, she reasoned that having some practice navigation became easier.

User *K* talked about the issue of paying attention: She reasoned that sometimes her attention was not allocated to the right label or button, which then influenced performance. To her it was not clear how to search a meeting since "I did not get that I should load of the available ones first". A consequence of this was that she had opened the meeting session

Table 3.4: Ease of Navigation Synopsis

1. Support users in focussing
2. In the navigation bar, only those features should be available that match the feature the user is currently running
3. Implement a looser organization of the meeting window

several times before she realized she should search for it from the search button.

What concerns the labels and buttons, participant *N* believed "they were standard for what he would expect, but not very innovative. I had to guess which ones did what." He reasoned that this was due to the fact that he was not a PC expert. However, his remark is quite important, as usability should support both experienced as well as inexperienced users.

Results From the Post-Test Questionnaire

Besides the comment section on navigation, all participants were also asked to rate the labels and commands as well as the ease of moving between different systems on a scale from 1 (worst) to 7 (best). We have already seen some examples.

Regarding labels and commands, the average of the IT-savvies, was at $4,85$ and the average of the IT-not savvies was at $5,2$, giving flexibility an overall rating of $5,0625$.

Regarding moving around windows, the average of the IT-savvies, was at $4,71$ and the average of the IT-not savvies was at $5,2$, giving flexibility an overall rating of $5,0$.

3.6.4 Navigation Synopsis

Ease of navigation was tested using examples where labels, buttons and windows were involved on the system part. On the user part, we have looked at the effects of paying attention to a task as well as the various levels of confusion involved due to the use of many different windows.

The users partly expressed liking the way navigation was facilitated. There is still room, however, for enhancing navigation (see Table 3.4 on page 47):

It is important to focus users in their work no matter what strategy they used. That means that only so much information should be highlighted or displayed that the user will need as a starting point for navigation.

In the navigation bar, only those features should be available that match the feature

the user is currently running. For instance, the user should not be allowed to mail any selection of the transcript before having loaded a meeting.

Furthermore, once the user has loaded a transcript, the numerous available windows may distract and confuse the user. Therefore, highlighting the transcript using increased saturation would enhance contrast between the items in that window.

Furthermore, the index tabs have been problematic in that they suggested users that they were headers for the discourse traces. A looser reorganization of that meeting window, especially the index tabs, would help the user focus and navigate more easily.

Process-based tasks tend to put the user into a sequential setup of tasks. Therefore, I observed that some users felt restricted in their eagerness to explore. This, however, is a completely different issue compared to system setup constraints, which forced the user to proceed against their own plan of action.

3.7 System Responsiveness And System Instructions

System responsiveness is an issue that can be associated with ease of user frustration. User frustration occurs when the user feels limited in his strategy, he is unable to solve a problem, or when the user has to wait a long time until he (maybe) receives feedback to his actions.

In this section, I will list problems and user comments that allow conclusions on system responsiveness as well as on system instructions in the Help/Documentation.

3.7.1 Tasks Testing This Issue

Main tasks involved for testing system responsiveness were Tasks 1, 6, 8 and 9. Task 1 provided insights on the loading times and system responsiveness regarding searches in the meeting explorer. Task 56 again tested this issue for defining meaning to the color lines displayed in the discourse feature window. I will extract information on system responsiveness from task 8 where users worked with the play feature. Finally, Task 9 will cover the aspect of the quality of system instructions.

3.7.2 User Strategies And Results

Loading the Meeting Browser

There were some interesting behaviors to observe even before the user actually began performing the tasks. Two types of users can be defined depending on how they dealt with loading the system.

The first group first read the initial task and then loaded the system. This group was dominated by IT-not savvies. The second group consisted mainly of IT-savvies. They opened the Meeting Browser first and while the system was loading they read the task description. Therefore, one can say that the IT-savvies made a time efficient start into the usability study.

Defining Meaning to the Color Lines

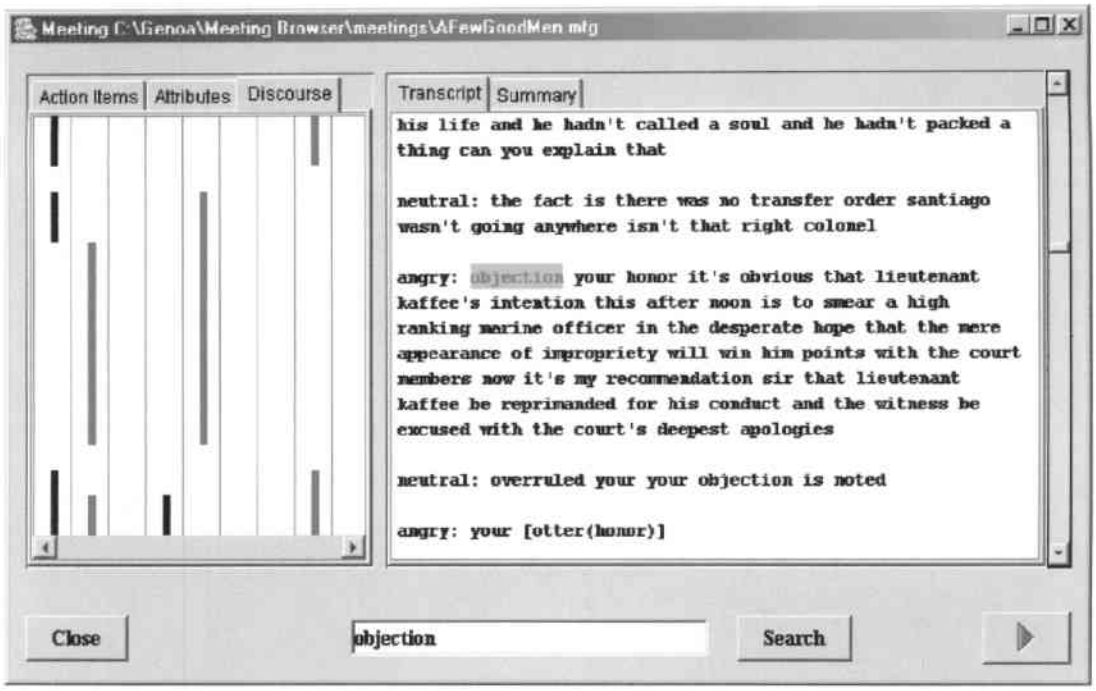

Figure 3.9: Within the transcript window, the user can find discourse feature in the left part of the screen. The color lines shown here, match the visual length of the paragraph currently displayed in the right part of the window.

During design of Task 6, I wanted users to slide over the color bars in order to extract the color meaning from the mouse-over. User *K* actually used that intended method and remarked that the system "displayed information about corresponding meetings whenever I browsed over the color strips".

Few users, however, had any idea of how to approach the lines. "What can I do with them?" many participants asked. So, some slid over the bars, but sometimes the labels did

either not show or the user moved on too quickly. Especially the IT-savvies were eager to find out more and directly clicked the colored bar, which can be attributed to their eagerness to explore. Some even played the paragraph before they were asked to do so in a later task.

Play Feature

As I said in a previous section, the users enjoyed using the play feature a lot. Many of them played the paragraphs over and over again in order to really grasp the contents. Having got the feel for the paragraph, though, after several plays, three participants wanted to stop playing. There was no response on part of the system, the paragraph continued to be played. Only one user succeeded in closing the play feature through its own file menu.

System Instructions

When participant *E* was trying to understand the properties window accessible through the color bars, he wanted to find out on what rating constraints the intensity of speech was based. When clicking help, however, he was unable to find information, as this part was not considered for the help feature yet.

3.7.3 Subject Comments

In the post-test questionnaire, users were asked whether they found the system responsive to their actions as well as how helpful the system was in coping with errors that were made.

Loading Responsiveness

In user *E*'s opinion, it takes too long to load and start the software.

Also, user *M* felt like loading took a short while, however, mentioned that it was okay then.

Play feature

Participant *C* described dealing with the play feature as follows: "I had first played a paragraph and then attempted to search for an item in the transcript. Then searching was not possible any more." He further mentioned that he had to choose big detours for achieving his purpose.

For user *J* is was a clear drawback of the system responsiveness that he was unable to stop playing a meeting excerpt when he asked it to do so.

General Observations

User *B* mentioned that the system was not responsive as it "did not recognize hitting the return key for starting a search".

User *D*, however, said that effects of actions could usually be immediately observed.

F concentrated on the search feature, which was not very powerful in his opinion.

Participant *N* discovered that the system hangs sometimes leaving him to guess whether his input had been accepted. He further reasons that it would be helpful if the system helped him in correcting his mistakes. He would see a solution in consulting the manual. User *P* would like to see error messages applicable for spelling correction in the transcript search. In her point of view, this would be especially helpful when talking about spelling corrections. Therefore, the messages could look as follows: "Did you mean 'friend'? You typed firend."

Results From the Post-Test Questionnaire

Besides the comment section on helpfulness in coping with errors, all participants were also asked to rate system responsiveness on a scale from 1 (worst) to 7(best).

Regarding helpfulness with errors that were made, the average of the IT-savvies, was at $3,6$, taking into account that users *C* and *F* did not rate this feature and therefore only five votes were counted. The average of the IT-not savvies was at $3,44$, giving system responsiveness an overall rating of $3,5$. This was the worst rating any issue received in the usability study.

Regarding system responsiveness, the average of the IT-savvies, was at $4,57$ and the average of the IT-not savvies was at $5,1$, giving system responsiveness an overall rating of $4,88$.

3.7.4 System Responsiveness Synopsis

This synopsis will summarize issues on system responsiveness and system instructions. Please £nd an overview on the results in Figure 3.5 on page 52 contains an overview of such recommendations.

M felt that the "system should be more aware of the user's intention and present solution alternatives". This idea would help many users in error detection. Oftentimes, if the system did not respond, the user was lost, not knowing how he should proceed.

If the system included some awareness function, this would enhance usability. In my point of view, such an awareness function should predict a user's intention based on previous actions done by the user. Eventually, the system would present the user with options on how he could proceed to achieve his purpose.

Table 3.5: System Responsiveness Synopsis

1. Present solution alternatives to the user performing a task through messages and through a thorough user handbook.
2. Include error messages with hints to the help feature.
3. Shorten loading times for the system
4. Allow to change the meeting window's size for getting a better view on the transcript.
5. When closing the meeting window all other windows associated with the transcript should also close.
6. In the help feature, focus on what the user needs to perform his work and get the search options to work.

This would be a long-term solution. On the short-term basis, error messages including hints would help. For instance, these messages could cover issues on how to proceed after applying a "Blank Search". Otherwise, when the user reaches the end of a transcript search, the user should be prompted whether he would like to continue from the beginning.

Additionally, a thorough user handbook designed after a minimalist approach would be necessary in the short-term. This handbook should both be implemented as an online and a print version such that parallel consultancy and testing is possible. Furthermore, this handbook should allow the users to explore and only cover the most important solutions for successful system use.

Chapter 4

Future Outlook And Feasibility

The post-test questionnaire was also used to find out about whether the system would be helpful in the user's job and what information the user would be interested in retrieving from the Multimodal Meeting Browser. Furthermore, this questionnaire aimed at disclosing future application areas for the system. I will present this information in the following sections.

4.1 Applicability of the System for the Participants

In the post-test questionnaire, the users were asked how useful they would rate the system to be in their job. The answers were manyfold.

One particular group of participants expressed, under which circumstances they would or would not use the Meeting Browser:

User *D* would use it, "if it were more refined". Right now, he feels limited by "current bugs in essential parts".

C believed that the system "is very helpful to create detailed transcripts". He, however, feels the need that important paragraphs should be marked according to the level of importance the user assigns it.

Participant *L* thought, "it would be useful if I had to prepare minutes on a weekly basis", and when he could get the system "for the right price".

At this point in time, user *M* reasoned, "there still are some problems with handling the program in terms of missing or not functioning buttons. Therefore, I would not use it when it comes to meetings with high-ranking personnel".

This group gave the system an overall rating of $5,75$ in usefulness for their job.

Another group mentioned what they would use the Meeting Browser for the following applications:

Participant *B* liked the idea that "I am able to check my meeting partner's utterances in hindsight for correctness".

E said: "I think that it would be very valuable, because I find myself often in the situation that I cannot recall exactly who said what. It is very helpful, since small comments might not find the way into the protocol, which might make a small difference for an individual".

G further mentions that "due to the large amount of meetings I have, it would be nice to have a reference to look up the topics discussed, especially if it is a long time ago".

Participant *K* appreciates the system as it is "useful for recording and summarizing that are in general very time consuming and distracting [activities] in meetings". The system to her was further useful for "other features such as replaying and summarizing".

This group on average rated the system's usefulness as 6.

The final group talked about the functionalities provided by the Meeting Browser that could be achieved with other tools also:

F "would save time searching through minutes, but this kind of technology already exists".

J explained that "I do not have to use meeting transcripts often enough to use the programme efficiently". The most useful tool for him was "the search function that I could use with Word also".

Being a legal clerk, user *L* gave the following reasoning: "because the meetings are generally only made up of myself and a client, which makes them reasonably simple, I do not need all the information the browser provides."

The overall rating given by this group was 4, 66.

The remaining users *A*, *H*, *I*, *O*, and *P* did not comment on the applicability of the Meeting Browser to their job. This group gave the system a rating of 4,2. Combining the ratings of all groups the system was regarded as useful in the participants' jobs at an average of 5, 13.

4.2 Information of Interest to the Participants

Question 13 asked the study participants: "When you analyze a meeting: What information would you be interested in?" The users had three general options, text information, participant information and emotions, which all are to some extent available in the current version of the Multimodal Meeting Browser. In category "Others" the users had the possibility to comment on more issues that are relevant to them for successful meeting

review. The following represents a selection of the results observed.

All participants agreed that text information provided by the transcript is the essential source of interest to the users. Participant *D* said that he would be interested in the full transcript with the ability to switch off unimportant parts right away.

Participant information seemed to be quite important for all except for two users, users *G* and *H*. Again, participant *D* had a more detailed idea of what participant information he would be interested in: names, contact info as well as a list of their contributions to the meeting.

User *H* demonstrated that she does not believe that emotions play a role for meeting or transcript analysis. As the color lines refer to the emotional aspect of the meeting, I will now give you an insight into the results: Of the 16 participants, 7 users, 3 IT-savvies and 4 IT-not savvies, rated the importance of emotions high. Some others, users *A*, *D*, and *L*, further pointed out that emotions were not important for them. They also remarked that such information is extremely difficult to capture due to the variety of emotional cues contained in speech and language.

Furthermore, the users shared other interests that could be considered for the Multimodal Meeting Browser:

Users *A*, *F*, and *G*, felt that they would like to directly be pointed to the conclusions or outcomes section of a meeting. It may be argued that such information is contained in a summary, however, in many structured meetings, important information is recaped in the end to ensure, for instance, that all tasks are assigned to the respective persons until the next meeting.

Timing and location seems to be important information to users *E*, *F*, and *O*. Participant *F* further suggested adding a direct link to connected meetings.

User *B* would like to see gesture and mimics analyzed for the emotional information. Related to the emotional information, participant *O* suggested an analysis concerning meeting dominance of meeting participants. User *C* finally suggested a glossary for the meeting transcript.

4.3 Participant Likes And Dislikes

Favorite Features "What features of the Multimodal Meeting Browser did you like best?" was asked in Question 10 of the post-test questionnaire. In the following I will present the results from this question. The results are also listed in Table 4.1 on page 56.

The Play feature was the favorite feature for 7 users, among them 4 IT-savvies, users *D*, *E*, *K*, and *P*, and 3 IT-not savvies, users *A*, *B*, and *J*.

User *D* pointed out that he liked "the ability to first read the transcript, and then also

Table 4.1: The Participants' Favorite Features

1. Play feature
2. Keyword search
3. Mail feature

hear the spoken words. Especially with the corresponding marking in the transcript, this allowed for easy grasping of the tone of the meeting".

Participant *E* explained this even further as "sometimes it is more important, how something has been said", and this is supported by listening to the voice playback. *J* remarked that he liked to be able to play the transcript without having to change the medium.

The second favorite feature to the users was the keyword search. 5 users liked this feature a lot. Among them 4 IT-not savvies, users *A*, *I*, *M* and *N*, and participant *P*, an IT-savvy user.

The third favorite feature was the mail feature that users *F*, *G* and *I*, 2 IT-savvies and 1 IT-not savvy, liked.

Furthermore, user *G* liked printing best, *H* liked having an accurate transcript, participant *M* liked voice recognition, and user *O* liked the structure of the system.

Least Favorite Features "What features of the Multimodal Meeting Browser did you like worst?" was asked as Question 11 in the post-test questionnaire. In the following, I will provide the analysis for this question. See also Table 4.2 on page 56.

Table 4.2: The Participants' Least Liked Features

1. Color representation in the discourse features
2. Layout, i.e. windows
3. Keyword search

About half of the users did not feel supported by the colors associated with the discourse features. 5 IT-not savvies, participants *A*, *H*, *I*, *L* and *O*, and 2 IT-savvies, participants *D* and *P*, expressed their dissatisfaction with the feature:

A remarked that "colors and colums were not explained in any part. So, maybe beams for the respective speakers could additionally be displayed in the colums. This could facilitate a visual search for the contributions of a certain meeting member." Also *D* does not understand the usefulness of the colour bars. Participant *I* said that the "task of colored bars was not immediately obvious, instead of colored bars, icons such as smilies may be useful".

User *L* remarked that the color lines "are not accurate enough, but they can not be, because the human voice is just too complex to be summed up in one color". User *O* shared the opinion that the discourse feature window needs to be more accurate and explaining, and *P* believed that the "discourse feature would need some differentiation on who said what".

There were two features that both were disliked by 3 users each. First, users *B*, *F* and *J* disliked the Meeting Browser's layout. On this topic user *B* said that "the meeting window is not visually attractive. It's not fun to use, but rather technical and bureaucratic". Participant *F* remarked that "there were too many windows" and "having to close one to get to another one is confusing". And user *J* further discussed that the system was "occasionally a bit IT-oriented related to things about remote servers, properties and so on".

Second, users *E*, *G*, and *N* did not like the keyword search. *E* mentioned that "the search function with the transcript was not obvious, that only one search can be performed and no other can follow, when one selects a certain paragraph and plays it following the first query". *G* was unhappy about not being able to search for issues related to the topic. Finally, user *N* remarked that he was only able to search for one keyword at a time.

It is interesting to see how diverse the perception on the keyword search is. It was mainly IT-not savvies who regarded that feature as their favorite one, while it was almost all IT-savvies who did not like the feature at all. User *G* gave her vote for the keyword search for both, disliked and favorite features, which is an unusual coincidence.

The fourth-worst liked feature was related to typesetting that was questioned by *K* and *P*. *K* said that "representing information in the meetings in plain style, i.e. capitalization issue, is not so suitable". And *P* added that "case sensitivity as well as spelling correction was not allowed".

Further problems concerned the following features:

On the topic of the unavailable right-click option, *E* said: "I would like to have right mouse button within the transcript window, which includes features such as email, print, copy and so on."

M did neither like the "Blank Search" nor printing, as he did not find a "Print Selection" feature and depended on help to return to the correct conceptual model.

User *O* did not like that it took the system so long to load and respond. Finally, *P* expressed her discontent about the fact that there was no button allowing to go back to the initial search window.

4.4 Suggestions for Future Applications

Question 14 asked the participants: "You have encountered meeting records of ordinary business meetings as well as court cases. What other application areas can you think of?" In the following, I will introduce the creative ideas that the users came up with. You can find a listing of the suggestions also in Table 4.3 on page 58.

Table 4.3: Possible Future Application Areas

1. University setting: lectures, exams.
2. Speech and discussion support: in the "Bundestag" and general assemblies.
3. Support for journalists
4. Support for psychological studies
5. Task scheduling
6. Sport team meetings
7. Spy test preparation or lie detection

Users *E*, *I*, *L*, and *O* were especially interested in the application of the Meeting Browser in the university setting, such as in lectures:

E would like to see a lecture application, as "this way students do not have to take notes, and teachers are able to recall what information has been given". Users *I* and *L* also suggests records of college classes, and *L* adds the possibility of taping oral exams.

Serving discussions and speeches would be an interesting application to users *C*, *E*, *G*, *L*, and *M*: as examples they name "Bundestag" discussions, general assemblies in companies as well as conferences.

Related to conferences, the Meeting Browser could also be used by journalists. This idea would be supported by users *E*, *J* and *K*. *E* believes that this would be "a help for journalists, this way they dont have to retype" the information. Furthermore, *K* mentioned "interviews and meetings concerning other issues such as economical and social ones

conveyed via TV". Related to TV broadcast, user *C* imagines using the Meeting Browser to extract close captions. User *M* would like to see political talk shows being taped and a transcript being extracted.

J sees a great potential for psychological studies, as "a friend of mine did analyze coaching sessions and had to type the text from videos into word to work with the text and create attributes, categories and so on".

Communication also plays a role for potential applications: user *D* believed voice mail transcription would be a good solution. *E* imagined that phone calls could be recorded and then analyzed. Also *I* suggested telephone conferences.

Furthermore, team-related applications would be possible:

P suggested "task scheduling and assignment, setting milestones especially relevant for the business or team meetings, something like Microsoft project".

Such a tool could also work for sport team meetings that user *A* talked about. In his point of view, such recordings "could be used as the basis for analysis after the match". Club activities were also the center of *G*'s attention. The Meeting Browser, in her point of view, could support celebrations and open days. Furthermore, she pictured medical information meetings including question and answer sessions as potential application areas.

Last but not least, *B* had the following suggestion: "I picture the Multimodal Meeting Browser on a palm top such that I can take it to contract meetings at a car dealer's. I would like to see a lie detector such that the car dealer cannot give me false promises. This could also be helpful during cross examinations." Also, user *N* saw a potential in spying test preparation.

4.5 Task Performance

In this final section, I would like to point out how the user performance rating can be viewed in the context of the user comments and their own rating of the features available in the current version.

As described in the methodology section, the calculation of the task performance looks as follows:

$$Task\ Performance = 1 - (1/\frac{task\ achievement}{task\ time}) \tag{4.1}$$

Task 1 was concerned with finding and loading a meeting. This task meant the first encounter with the system. It is necessary to emphasize the connection between the results of both task achievement and task time. A good performance does not imply that it also was a quick performance, and vice versa. This could also be observed in this task.

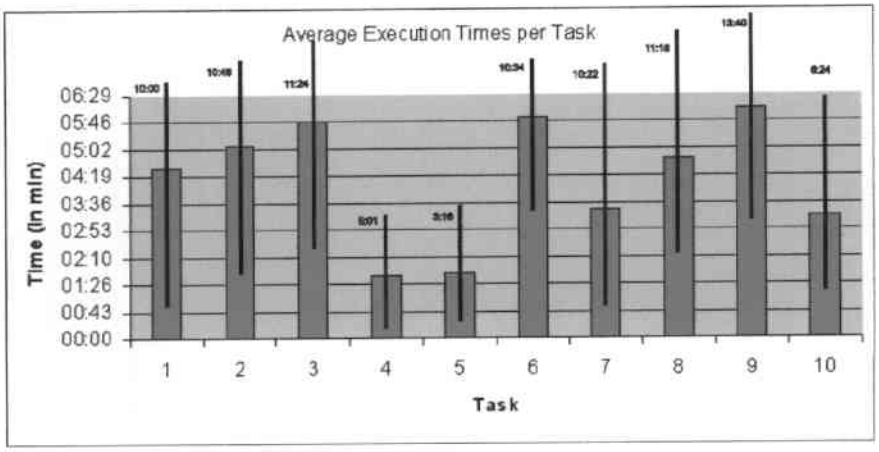

Figure 4.1: This graph presents an overview of average, minimum and maximum execution times that users needed for the respective tasks. The average is generally displayed as the red-tinted bars, while minimum and maximum times are indicated by the lower and upper end of the black lines.

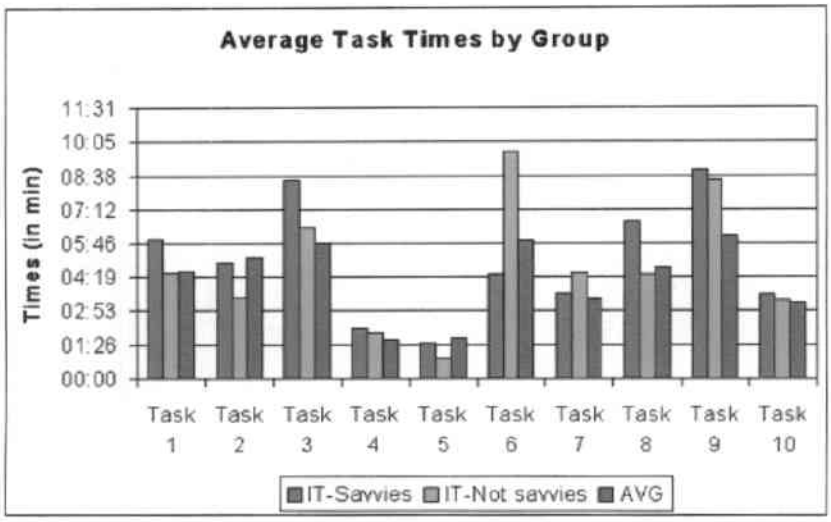

Figure 4.2: This graph presents an overview of averages of task execution times, organized respective to three groups: the IT-savvies, the IT-not savvies as well as the average of all participants.

The performance ranged from $60,00$, user *H*, to $98,51$, user *I*. This number could be connected to a very quick performance of user *I*, which was 1.48 minutes as well as a very slow performance by user *H*, which was 10.00 minutes. Also their task achievement was reflected in their performance rating as *H* received 25% and *I* received 100%.

The average task performance was $92,39$ and the average time needed to perform the

task was 4.33 minutes, which meant that Task 1 was seventh out of ten tasks comparing the ranking of average performance.

Task 2 concerned searching for keywords. The performance ranged from 86, 05, user *B*, to 98, 09, user *M*. This number could be connected to a very quick performance of user *M*, which was 1.43 minutes as well as a very slow performance by user *B*, which was 10.46 minutes. It is interesting to note, that both of them were not error free in performing the task, and that they therefore received 75%. Others that performed 100%, such as users *A*, *F* and *O*, had very good performance ratings, however as their task times were higher compared to *M*, the latter achieved an result.

The average task performance was at 92, 86 and the average time needed to perform the task was 5.08 minutes, which meant that task two was sixth out of ten tasks comparing the ranking of average performance.

Tasks 1 and 2 both were concerned with searching, either searching for a meeting or searching within a meeting. Even though many users took a while to succeed in these dependent tasks, they were quite impressed by the search functionality. This was mirrored in the rating for searching meetings asked for in post-test Question 9: They gave this function the best rating for any feature in the Multimodal Meeting Browser. I track this back to the Meeting Browser's potential in improvement as well as the users' forgiveness concerning the "Blank Search" problem.

Task 3 concerned mailing a selection. This task demonstrated the worst overall task performance in the user test, which was at 85, 24. This results reflects the problems nearly all participants encountered when trying to find the "Mail Selection" option in the system. Those users had an advantage that worked menu-based: User *I* chose this strategy and performed the task in 2.13 minutes with a 100% task achievement and had the performance rating of 97, 83. User *C*, however, had more than six deviations, meaning 25% task achievement, and needed 11.24 minutes to succeed in that task. This left him with only 54, 95.

Even though average task performance was worst, the average task time was at 5.46 minutes. Knowing about the constraints tells us that there must have been a lot of deviations. In fact, nine participants had five or more deviations.

Task 4 concerned loading another meeting. Taking into consideration that this was quite a short task and that this was the second attempt to load a meeting, the users had a very successful performance. In fact, this task had an average performance rating of 98, 40 as well as an average time of 1.39 minutes. The range of this task performance was from 94, 98, user *L*, and 99, 90 for user *G*. *G* furthermore only took 0.10 minutes to perform the task and *L* needed 5.01 minutes.

Task 5 asked the users again to look for keywords. This task was also taken into

consideration for learning effects. The users learnt from their experience, which was also reflected in the task performance. This task was second in the average task performance with 98, 01. This again can be attributed to very quick task execution times with averaging at 1.43 minutes. The performance rating ranged from 95, 54, user *B*, and 99, 75, user *O*.

Task 6 asked the users to assign meaning to the color lines in the discourse features. This task was conceptually difficult and therefore required active thinking and reasoning by the participants. Along with Task 1, this task was seventh having an average of 92,39. This number, however, reflects less deviations and longer execution times on part of the user. The average execution time in this task was 5.53 minutes compared to 4.33 minutes in task one. The performance rating ranged from 86, 20, user *A*, and 96, 92, user *J*.

Task 7 asked the users to examine the properties window associated with the discourse feature. The average task performance is mainly based on the good task time by the users, which was 3.25 minutes. It led to an average task performance of 95, 53, which made this task third in the listing. Participant *H* had the worst performance with 86, 78, and user *M* the best with 98, 97.

Task 8 allowed the user to play a passage from the meeting transcript. For most users this was fun to do and many performed very well. The only problem was related to time in this task pulling the average task performance down to 94, 73. User *E* was best with 98, 42 and user *C* had the worst performance with 85, 07. On average, the participants needed 4.46 minutes to perform the task.

Task 9 was also ninth in the ranking of the average performance. The participants were asked to consult the help feature. The average performance rating was 91, 75 and can be attributed to time as the average task time was at 6.06 minutes. User *C* performed worst with 75, 66 compared to user *O* who received 96, 86.

Task 10 was the last task asking the participants to print a selection of the transcript. Average task performance was 95, 30. User *K* had the lowest rating with 87, 51 and user *G* the best with 98, 80. The average task time was 3.17 minutes and also the reason why the task was fourth in comparison to the other nine tasks.

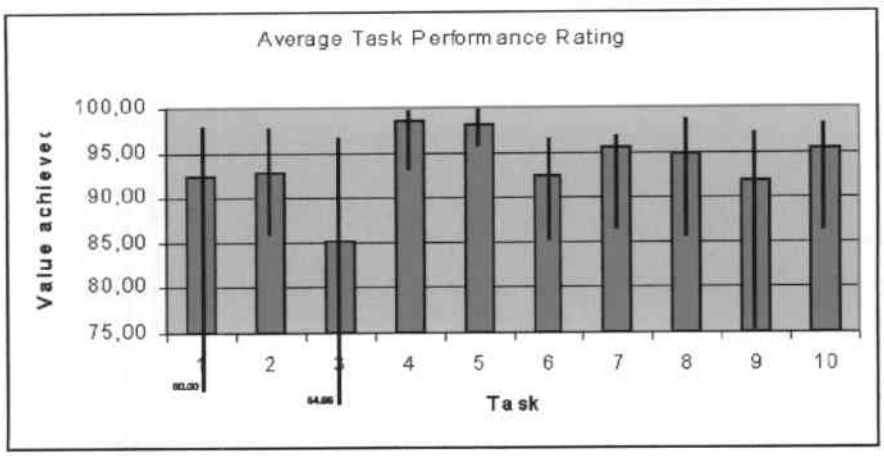

Figure 4.3: This graph presents an overview of average, minimum and maximum task performance that users achieved for each task. The average is generally displayed as the red-tinted bars, while minimum and maximum times are indicated by the lower and upper end of the black lines.

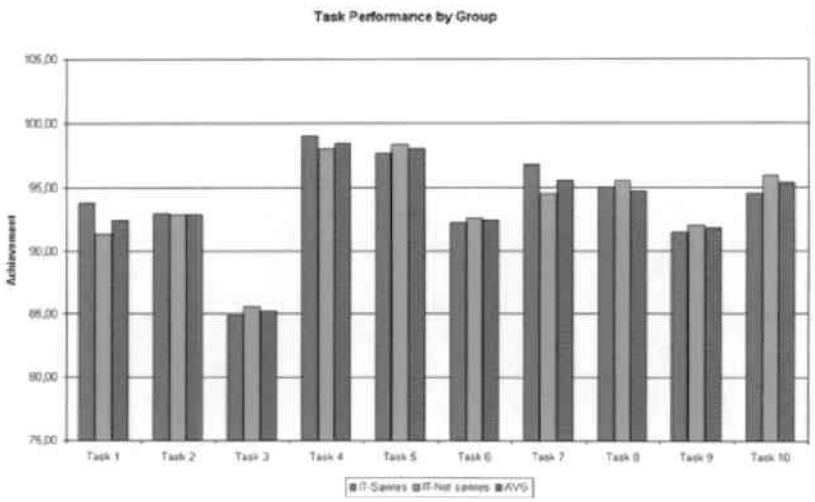

Figure 4.4: This graph presents an overview of averages of task performance, organized respective to three groups: the IT-savvies, the IT-not savvies as well as the average of all participants.

Chapter 5

Conclusion

This final chapter talks about general observations and conclusions that can be drawn from the usability study results. Reviewing meeting records has proven to be an activity that involves manyfold features, difficulties as well as opportunities. In the beginning, I introduced four main questions that guided me through this analysis.

Acceptance of the Meeting Browser by the Users The users were quite intrigued by the system. They explored and examined the system and they were eager to work with it. Playing transcript paragraphs was one of their favorite activities.

Nonetheless, they often felt some limitations due to system bugs, errors and unknown features. The participants struggled with features and it sometimes made them frustrated when they were unable to move ahead.

Still, user acceptance was quite good. During the post-test questionnaire session, the users in general gave all system features very good ratings on a scale from 1, worst rating, to 7, best rating. This attitude can mainly be associated with the high potential that participants saw for the system. They did not always judge that potential for their own purposes, but for the purpose of efficient lectures, for instance.

User Strategies Used for Successful Task Completion Bhavani and John (2000) already considered arguments related to efficient strategies. These include, first, being "efficient because they exploit powers offered by computers such as iteration propagation, and visualization". Second, these strategies "need to be made explicit to users because the knowledge to use them is suggested neither by tools nor by task descriptions". Their final argument for strategies is that "they are generally useful across computer applications".

Related to these arguments, the Multimodal Meeting Browser shows potential in that the system can be extended such that the right-click option is available at all times and includes links to features that are already available in different locations in the system.

Another opportunity in this case is that the right-click options can be made dependent on current system states, i.e. whether a selection has been made or not.

At this point in time, only few user strategies and conceptual models are supported by the Meeting Browser. There were quite a few participants, however, that were unable to apply their strategies due to system limitations. Some users therefore felt forced into one certain path to achieve a task.

Problems Encountered During Task Execution The problems users faced while performing a certain assignment included the following:

Besides feeling restricted in their strategies, the participants encountered problems with the search features: It is clear to see that the "Blank Search" needs to be eliminated. This, however, was not the only search problem. The conceptual model, referring to the logic underlying the search interface, was not clear to all users and using Bhavani and John's (2000) words, "prior knowledge dominated performance".

Furthermore, the keyword search made users feel limited in their actions as they were only able to search for one keyword, there was no spelling correction, and once you have navigated the list of results, you were unable to return to the first result.

Users also encountered difficulties when they did not find a mail menu item in the navigation bar and were misled by the icon "Mail Transcript". Furthermore, the existence of the "Mail Selection" as a direct access option, suggested users that there was also a "Print Selection" option available. This distracted them so much that they were unable to recall the generally known procedure.

The feature that involved the most conceptual mistakes was the discourse feature window and its color lines. The problems in this case were not limited to what was hidden behind the lines, but especially what the colors represented. As discussed in the relevant sections, the problems associated with the color lines were complex ones, and therefore require careful rethinking, change of layout or, at least, a thorough description and guidelines in the help feature.

Opportunities and Future Outlook for System Success As mentioned earlier, some users would like to see a Meeting Browser application for lectures in order to facilitate their own understanding, to aid preparation for the next class, to be able to recap what the instructor taught at a later instant, or to be able to get precise information of a lecture's content after being sick.

According to 2 users, the Meeting Browser could serve as a lie detecting system and therefore as proof for promises that were not met by a car dealer, for instance.

Most users saw the greatest potential for the Meeting Browser concerning applications

areas such as conferences of any sort or for communication purposes.

In fact, 11 users often use instant messenging systems like AIM, ICQ and/or MSN Messenger. These users emphasized the applicability of communication support such as voice mail transcription, phone calls or telephone conferences.

Based on this information, one may consider an online implementation of the tool. Allowing this, however, would raise the question of how to incorporate all the different windows in one single web browser window?

This leads to the next point: As the users encountered difficulties working with all the windows in the Meeting Browser, designing a new layout of the tool including all features in one big window, could be used as an intermediate step of making the Meeting Browser a web-based application.

Another great opportunity lies in the help system for the Meeting Browser. Redesigning the help feature could be achieved by using Carroll's (1997) minimalist approach and taking into consideration user behavior patterns and the tradeoffs going along with them.

Concluding Observations In the end, I would like to name two general observations related to group specific behavior:

The IT-savvy participants generally critiqued less issues on the system than did the IT-Not savvies. When asking them about flaws in the system, they often mentioned that "you have to get used to it" and then working around problems would be okay. I argue that these users are working with lots of imperfect systems. Therefore, many IT users do no longer appreciate or do not desire using a perfectly usable system.

Another general attitude that I observed was the following: The users agreed in the attitude that they would rather adapt to a system's functionalities to achieve a certain goal, instead of requiring the system to be most suitable and usable for them.

This attitude goes hand in hand with the first observation, but somewhat clashes with it as well: On the one hand, users argue that they felt limited in using their own approach for reviewing the meeting records, on the other hand, they would adapt to the system setup for better performance.

Testing the Multimodal Meeting Browser's usability for reviewing meeting records in an empirical study setting like this one, was the first step for the system to become even more efficient and valuable to the users. Using this version already filled users with excitement. This was especially reflected in their ratings and comments.

This thesis aimed at capturing the general attitude of users towards the system, revealing its potential for improvement as well as giving a future outlook and recommendations for future implementations. An important factor considered throughout the study and its analysis were user strategies. These need to be considered further to support user action

and make the system usable, fun to use as well as visually and conceptually attractive in the future.

Bibliography

Bett, M., Gross, R., Yu, H., Zhu, X., Pan, Y., Yang, J. & Waibel, A. (2000). Multimodal Meeting Tracker. In *Proceedings of RIAO 2000*. Paris, France

Bhavani, S. K. & John, B. E. (2000). The Strategic Use of Complex Computer Systems. In *Human-Computer Interaction* (Vol. 15, pp. 107-137). Lawrence Erlbaum Assiciates, Inc.

Brockmann, R. J. (1990). The Why, Where and How of Minimalism. *ACM SIGDOC Asterisk Journal of Computer Documentation, Proceedings of the conference on SIGDDOC'90*, 14, 4, 111-119.

Carroll, J. M. (1997). *Reconstructing Minimalism*. Cambridge, MA: MIT Press.

Hirst, S. J. (1995). *HyperLib Deliverable 5.2: Supplementary Report on the Evaluation of the Hyperlib Interfaces*. Retrieved August 20, 2002, from http://143.169.20.1/MAN/WP52/root.html.

Mayhew, D. J. (1999). *The usability engineering lifecycle: a practitioner's handbook for user interface design*. San Francisco: Morgan Kaufmann Publishers, Inc..

Suh, B., Woodruff, A., Rosenholtz, R. & Glass, A. (2002). Popout Prism: Adding Perceptual Principles to Overview+Detail Document Interfaces. In *Proceedings of SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves* (pp. 251-258). Minneapolis, MT

Whiteside, J., Bennett, J. & Holtzblatt, A. (1988). Usability engineering: Our experience and evolution. In M. Helander Ed., *The Handbook of Human-Computer Interaction* (pp. 791-817). Amsterdam: North-Holland.

Wixon, D. & Wilson, C. (1988). The Usability Engineering Framework for Product Design and Evaluation. In M. Helander, T. K. Landauer, & P. Prabhu (Eds.), *The Handbook of Human-Computer Interaction* (2nd ed., pp. 653-688). Amsterdam: Elsevier Science B.V..

Appendix A

Task execution times listed by user

User \ Task	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	Total
A	07:00	04:00	05:48	02:42	01:20	10:34	04:42	04:04	13:40	03:00	**56:50**
B	02:38	10:46	02:42	01:15	03:34	07:06	04:29	02:14	05:19	01:18	**41:21**
C	08:00	06:27	11:24	03:18	01:24	05:18	02:41	11:18	12:14	04:02	**06:06**
D	04:33	05:19	02:48	00:28	00:26	04:30	01:27	02:20	02:38	02:57	**27:26**
E	03:56	06:14	10:03	01:30	01:04	03:42	03:27	01:57	04:56	03:39	**40:28**
F	04:28	02:40	04:17	00:30	02:29	09:01	02:40	05:59	04:24	01:30	**37:58**
G	05:24	06:56	05:33	00:10	03:16	05:22	01:32	02:49	05:35	01:19	**37:56**
H	10:00	06:28	06:22	02:30	01:22	05:58	03:30	05:28	05:06	02:48	**49:32**
I	01:48	03:25	02:16	01:02	01:24	05:40	03:42	05:05	04:07	03:36	**32:05**
J	05:00	06:38	05:30	01:31	02:46	04:25	10:22	07:07	09:39	03:46	**56:44**
K	02:04	05:40	07:04	01:02	03:17	07:50	01:07	06:07	04:00	06:24	**44:35**
L	03:44	03:17	04:11	05:01	00:38	04:00	03:36	10:02	08:54	04:24	**47:47**
M	04:18	01:43	04:16	01:08	00:55	04:24	01:02	02:26	04:54	04:06	**29:12**
N	04:00	06:18	07:24	02:06	01:42	03:57	01:31	02:24	03:22	02:50	**35:34**
O	02:00	02:56	07:00	01:11	00:24	08:40	04:16	04:51	03:13	03:40	**38:11**
P	03:49	03:29	05:34	01:02	01:33	03:38	04:37	02:06	05:31	03:06	**34:25**
AVG	**04:33**	**05:08**	**05:46**	**01:39**	**01:43**	**05:53**	**03:25**	**04:46**	**06:06**	**03:17**	
										Total	**00:42:16**

Figure 1: This figure lists the execution times that the user needed for performing a certain task.

Appendix B

User	1 Usefulness	2 Labels & Comands	3 Layout & Presentation	4 Ease of navigation	5 Flexibility in Use	6 Helpfulness with errors	7 Responsive-ness	8 Mail a Selection	9 Search Meetings	AVG
A	3	5	5	6	2	3	6	5	6	4,555556
B	5	5	5	5	3	3	2	6	6	4,444444
C	6	5	5	6	3		2	4	6	5
D	6	5	6	7	4	4	6	6	5	5,444444
E	7	5	5	5	6	5	3	7	7	5,555556
F	6	6	1	3	2		5	7	4	4,25
G	6	4	6	3	4	3	5	6,5	4	4,611111
H	4	5	3	4	5	4	5	5	4	4,333333
I	5	7	5	5	6	5	6	7	7	5,888889
J	4	5	6	6	6	4	4	7	6	5,333333
K	6	4	3	4	7	3	5	6	4	4,666667
L	4	6	6	5	7	2	7	6	7	5,555556
M	5	6	5	5	3	2	5	6	5	4,666667
N	6	4	4	6	5	4	6	7	6	5,333333
O	4	4	5	5	5	4	5	6	5	4,777778
P	5	5	4	5	4	3	6	6	6	4,888889
	5,125	5,0625	4,625	5	4,5	3,4	4,875	6,09375	5,5	4,956597

Figure 2: In the post-test questionnaire, the users indirectly rated the testing issues of the system. This figure includes the ratings given by each user for every issue.

Appendix C

Questions of the Post-Test Questionnaire

The following table contains the list of questions asked all participants in the post-test questionnaire. Question 1 through 10 dealt included ratings from 1 (worst) to 7 (best) as well as an explanation part. The remaining 5 questions were verbatim questions only.

Question 1	How useful would you rate the system to be in your job?
Question 2	Were the labels, options and commands on using the system easy to follow and use?
Question 3	Was the layout and presentation of information clear?
Question 4	Was it easy to move around the different parts of the system?
Question 5	Did the system give you sufficient □exibility to work in the way you wanted?
Question 6	Was the system helpful in coping with any errors that were made.
Question 7	Did you find the system responsive to your inputs?
Question 8	Do you consider the system suitable for the task of extracting meeting paragraphs and mailing them to other participants?
Question 9	Do you consider the system suitable for the task of finding meetings?
Question 10	What features of the Multimodal Meeting Browser did you like best?
Question 11	What features of the Multimodal Meeting Browser did you like worst?
Question 12	How did you feel about dealing with the different windows involved in the Multimodal Meeting Browser?
Question 13	When you analyze a meeting: What information would you be interested in? Text information, Participants, Emotions by participants and/or Other?
Question 14	You have encountered meeting records of ordinary business meetings as well as court cases. What other application areas can you think of?