

Building a Digital Library of Captured Educational Experiences

Gregory D. Abowd, Lonnie D. Harvel and Jason A. Brotherton
{abowd, ldh, brothert}@cc.gatech.edu

College of Computing & GVU Center
Georgia Institute of Technology
Atlanta, Georgia 30332-0280 USA

Abstract

Since 1995, we have been researching the application of ubiquitous computing technology to support the automated capture of live university lectures so that students and teachers may later access them. With virtually no additional effort beyond that which lecturers already expend on preparing and delivering a lecture, we are able to create a repository, or digital library, of rich educational experiences that is constantly growing. The resulting archive includes a heterogeneous mix of materials presented in lectures. In this paper, we discuss access issues for this digital library that cover short-term and long-term use of the repository. While our efforts have focussed quite a bit on short-term access that occurs over the duration of a course, it is clear that significant value is added to the archive as it is tuned for long-term use. These long-term access issues for an experiential digital library have not yet been addressed, and we highlight some of those challenges in this paper.

1. Introduction

One possible service in future computer environments is the capture of everyday experiences, making that record available for later use. A general challenge in ubiquitous computing is to provide automated tools that support the capture, integration, and access of this multimedia record [5]. By providing automatic capture, we free humans to do what they do well: attend to, synthesize, and understand what is happening around them. One view of automated capture is as a paradigm for multimedia authoring [20].

The resulting archive is a digital library of capture experiences that has many potential uses and presents a number of research challenges. Of the many definitions of digital libraries the one we apply here is: *a collection of digital objects, including text, video, and audio, along with methods for access and retrieval, and for selection, organization, and maintenance of the collection* [7]. In this paper we discuss the capture, storage, and access of

live lectures, and the digital library related issues of the work.

The eClass project¹ is an attempt to support both teaching and learning in a university through the introduction of automated support for lecture capture. In our approach a room has been enhanced with the capability to record a lecture experience. Throughout the project, we have gained extensive experience in the use of an instrumented environment to support automated capture. One of the main goals of the project was to build an instrumented classroom environment in which capture was reliable and easy. We have produced a system that has captured parts or all of over 2900 lectures at the Georgia Institute of Technology. Courses have ranged from graduate to undergraduate and have been in areas of computer science, electrical engineering, and mathematics. The eClass is currently in use at other universities as well. This continued interest and use of the system both locally and elsewhere is one clear measure of success.

1.1. Overview of paper

We begin in Section 2 with a brief description of the structure of the eClass system that facilitates automated lecture capture and access. The resulting archive clearly fits the definition of a digital library, though the nature of the archive poses problems of its own. The archive is organized as a collection of courses. Courses are a collection of lectures, represented as hyper-documents. An analogy to traditional education archives is to equate courses with books, and lectures with chapters in those books. As we will discuss, the segmentation of the lectures into “paragraphs”, or conceptual blocks based on topic, is one of the many challenges the archive presents as a digital library.

2. The eClass Infrastructure

The eClass project exploits ubiquitous computing technology, in an instrumented space, to capture a varied

¹ See <http://www.cc.gatech.edu/fce/eclass> for more information. This project was formerly called Classroom 2000.

set of interactions between the instructor and the environment. The aim of the project is to allow students to take on a more active and enriching role in the classroom by capturing details of a live session and making them widely available. The student accessibility of the captured lecture experience reduces the need for copying material; allowing for more interaction and attention to commentary and annotation in the student record.

The eClass system is structured in four phases — pre-production, live capture, postproduction/storage, and access [3].

Pre-production: We made a conscious decision to minimize the work required by an instructor in this phase. It is possible to use the system without any preliminary effort. However, many instructors produce graphic slides in advance of the lecture.

Live Capture: Once a lecture begins, we attempt to capture, with a time stamp, all of the relevant activity. We understand the live experience as a generator of various streams of information. For example, the set of prepared slides that might be shown during the lecture constitutes one stream that is divided into discrete chunks, or slides, that appear at various times. Handwritten notes, made either by the teacher or a student, constitute another stream, with individual strokes (pen-down to pen-up) defining the time-stamped chunks. Other examples of streams include the audio or video recorded by one or more microphones or cameras, a series of URLs (uniform resource locators) visited with a Web browser during a recorded session, or even a series of user interactions with a program that is being demonstrated during the lecture. This phase ends when the live session, such as the class lecture, ends.

The eClass software has been used to support activities in several classrooms and meeting rooms with varied hardware infrastructures. The most comprehensive hardware installation is found in the original prototype classroom, depicted in Figure 1, which has been in operation since January 1997. In this room, the hardware infrastructure includes an electronic whiteboard, two ceiling-mounted projectors, one camera and several microphones embedded in the ceiling, all operating in unison through a suite of client-server software modules.




Figure 1: The Classroom 2000 classroom is instrumented to capture of the information presented on the electronic whiteboard (A), captured by a video-camera in the back of the room (B) and by microphones embedded in the ceiling (C), as well as web browsing projected (D). Information from previous slides is presented simultaneously in other projected screen (E)

The supporting software, referred to as the Zen* system, performs the tasks of capturing and synchronizing streams of information during each live session. These tasks include controlling the capture of the media streams (audio and video) and orchestrating the generation of an associated document during postproduction.

Postproduction/Storage: The whole purpose of defining and capturing streams is to support the development of interfaces that integrate or associate related streams. For example, we wanted to be able to look at any part of a lecture after class and associate it with what was being said during that part of the lecture. This integration activity is central to the postproduction activity. We developed a technique for post-processing captured lecture material that could facilitate later browsing [11]. We currently place most of the post-processed material in a database and file system that facilitates a variety of interfaces to browse the captured material.

Access: The final phase involves the student (and teacher) access of some portion of the live lecture experience. The interfaces used by the students are very different from the physical interfaces used during live capture. For example, very few students would be expected to possess a networked, pen-based computer, so we cannot assume that form of interaction will be available. Since access needs to be universal and across platforms, it is important to build on an existing infrastructure, such as the Web. In the next two sections, we describe some of the interfaces for accessing lecture experiences. A more complete description of these interfaces can be found elsewhere [2], [20].

3. Short-term Access

The document presented in background of Figure 2 is a sample of the hyper-documents automatically generated with the eClass infrastructure. In this case, the frame on the right hand-side presents the lecture as a sequence of discrete slides as presented or created during the lecture. The frame on the left presents a timeline decorated with relevant activity as it happened during the lecture, either a new slide being visited on the electronic whiteboard or a Web page being visited on a separate display inside the classroom. The frame on the bottom left of the window presents the video (and audio) recording of the lecture, played back as a RealNetworks™ stream. The window on the foreground corresponds to a web page visited during the lecture — it was open by selecting its title used to decorate the timeline.




Figure 2: Web browsers are used for presenting the captured notes that link presented material from the electronic whiteboard and Web pages visited in class with streaming digital audio/video recorded during class

The eClass system contains a variety of sources including: slides, captured ink, audio, video, URLs, transcriptions of handwritten information, transcriptions of voice, titles of web pages that have visited during a lecture, and annotation that is added by the instructor after a lecture.

Viewing a single lecture is a very short-term access requirement. As a term progresses and more lectures are captured, it is often useful to be able to browse the entire contents of a course. We use an automatically generated syllabus of lectures to facilitate this form of browsing. Lectures are arranged in reverse chronological order.

Since the captured lecture information has been loaded into a database that contains all of the information from all other courses, a user can initiate a search over

that lecture, or over all lectures for the course. It is also possible to look for information in all courses or only in selected courses using the interface shown in Figure 3.



Figure 3: Search over several courses

An alternative "advanced" interface as shown in Figure 4 allows the user to select the lectures and content upon which the search operations are to take effect, which content to consider, and to specify whether or not the search should be case-sensitive.



Figure 4: Advanced search interface

Given a query, the search engine searches the database for the specific terms. Figure 5 shows an example of a screenshot that results from searching for "Spiral Model" over a subset courses in the database. The figure shows the detailed view of the search results that is available for each lecture. Thumbnail images are shown along with a list of where each occurrence appears. These images serve as anchors that can be selected by the user to view the lecture in a separate window, automatically keyed to that specific slide.




Figure 5: Detailed view of a search for “spiral model”.

4. Long-term Access

A typical course at Georgia Tech runs for fifteen weeks, and consists of 45 lecture hours. This amount of captured information can be difficult to search and manage. This situation is compounded by the fact that students often take several courses each term and many more courses over their entire academic experience. The capability of searching the content of many different courses may help a student find a relevant piece of information, or see how different courses are related along a common topic.

We have not thoroughly investigated the variety of interfaces and retrieval techniques to facilitate this long-term access task. In order to augment a users ability to browse a large collection of captured lectures, we prototyped a new interaction and visualization technique, the multi-scale timeline slider, or MTS [23]. The MTS allows a user to navigate a stream of information and events as well as control the playback of that stream. User controlled, multi-scale zoom is supported through a set of individual timelines that focus the stream’s context. The user interacts with the timelines by creating focus regions, manipulating existing focus regions, and manipulating the playback control.

The visualization consists of a series of individual timelines, where each subsequent timeline is a focused region of the previous timeline. Each timeline represents a series of ordered streams. The recessed gray rectangles on the timeline represent a stream. Each stream is annotated with events, drawn as colored lines and boxes above and below the stream. In Figure 6, the streams are videos of meetings and the events are where different

keywords, such as system components, were discussed. Information is not added as the timeline becomes more focused. However, as focus increases, the individual streams and events spread out and are easier to distinguish.




Figure 6: The Multi-Scale Timeline Slider interface.

5. Digital Library Challenges

We have demonstrated through the eClass system that it is possible to overcome the difficulties of authoring rich multimedia documents that support access to live educational experiences. Though we have investigated some of the long-term implications of a digital library of these experiences [8], many challenges remain. We outline some of these in this section.

5.1. Indexing & Weighting

The contextual scope of a search ranges from a single course to several courses. Since eClass records the time of capture, searches may also be conducted over an extended period of time.

The existing search system is based on SQL queries of the database without the use of indices, weighted or otherwise [20]. The use of indices would improve the ability of the system to provide relevance searches. However, the structure of the document creates difficulties with traditional indexing systems.

The eClass hyper-document is not only a multi-component artifact; it has the added dimension of time as a contextual reference to the material. In order to correctly index and weight a term, the system must know in what component a term occurs, to what hyper-document the component belongs, and at what time within the stream the term occurs. If the granularity of the weighted index is on the lecture level, the information can be acquired easily from the eClass database. However, if the intent is to provide a smaller granularity, problems arise in aligning the individual components temporally.

The collection is divided in two different ways, using time. *Horizontal Time* is used to align separate components of an individual lecture. *Vertical Time* is used to align individual lecture component or segments with past or future components. The dual aspect of time within the collection affects indexing and weighting as well as topical analysis and summarization.

As an example of horizontal time, we consider the lectures that contain voice transcripts. The transcript time stamps do not align uniformly with the access times for individual slides. One possible method is to use the additional context of the ink strokes to assist in aligning the transcript with the slide. However, the ink strokes reference the time index of the audio stream, and not that of the transcript. Though the two are close, they are not identical. A second approach would be to use a topical analysis of the disparate components to provide an alignment based on the topical structure as well as the temporal. This addresses the question of what to do when the captured audio is not actually relevant to the slides that occur within the same temporal slice. (that is, an instructor engages in an extended side-track from the original topic expressed on a slide.) This approach requires the existence of a topical analysis, which will require keyword selection and summarization.

The second division of the collection is based on vertical time. One approach, built upon the semi-structured nature of lecture material, is to use recurrence in time as an element in the calculation of a compound weight for a term [22]. Previously we discussed using topical analysis to align contemporary streams along the time dimension. It is also possible, assuming dependable keyword selection, to align segments of related material over time. Other contextual clues can be used to assist or even guide the process of vertical alignment. In the case of eClass, the recurrence of courses can provide a rough-pass alignment of the collection.

For example, CS6750a is taught in the Fall and then again in the Spring. Likewise, it has been taught in previous years. This information can be used to align the collection on the course level. Using information collected explicitly from the slides, some alignment can be gained on the lecture level. However, to accurately align lectures, or more likely, segments of lectures, analysis of information implicit in the collection must be attained. Once a lecture, or lecture segment, has been aligned with compatriot material from previous and future experiences, an attempt can be made to identify predominant keywords within the component over time. This can provide stronger weights for material that is covered intentionally versus those that are incidental or only supportive of the planned material [26].

There are, of course, other issues related to the weighting of terms within multi-component, time-

sensitive collections. Should a term found in the prepared slides be weighted equally with a term found in the voice transcript? Exactly how should time affect the weighting of terms? Should terms within more recent document components receive higher relevance weights than those from previous semesters? Since course material is routinely reused, the question of weighting over time is relevant. Likewise, since faculty sometimes shares course materials, is the authorship of one component relevant to the weighting of that component within a hyper-document? Is it possible to create a weighting methodology that is independent of the lecturer, or will the weighting depend upon the presentation style?

As mentioned earlier, the eClass system also allows for the archiving of referenced URLs. Web pages that are visited during the captured lecture are referenced within the timeline of the hyper-document. This poses another challenge to conventional indexing and weighting. Should terms found within a referenced website be included in the index of terms associated with a hyper-document in the archive? If they are included, how should they be weighted? A possible approach is to allow terms within the web page to effect the weights of existing terms in the hyper-document without increasing the set of terms established by the local components.

Finally, since eClass is a heterogeneous collection of streams, how do we accurately compare lectures or courses, which have no transcripts, with those that have a wider collection of material? This problem is similar to that of comparing indices based on whole documents with those based on abstracts only.

5.2. Summarization

The eClass system, like other hypermedia archives, poses challenges to the existing summarization technologies. The multi-component nature of hyper-documents within the archive is similar to the problem of multi-document summarization [24], [16]. Summarization of courses would pose a similar problem, requiring the inclusion of multiple lectures. The archive is also a multimedia collection, requiring the use of new approaches in multimedia summarization. Since the hyper-documents in the eClass archive contain text components as well as media components, topical analysis applied to the text components can be used to augment the summarization of the media components, as in the work done by Takeshita [27].

As a foundation to both summarization and indexing, topical analysis of the hyper-document is an important aspect of making the eClass archive more useable as a library of captured educational experiences. There are many challenges to topical analysis, which we are beginning to explore.

As with indexing and weighting, the eClass archive adds time to the mix, creating new challenges for summarization as well. Summarization of material through time, and preserving the sense of time, is essential in an experiential archive.

Existing methods must be explored, and new methods created to segment captured data into conceptual blocks. Data segmentation routines and key term detection strategies need to be developed to account for both the multi-component nature of the hyper-document as well as the temporal nature of the data stream. The concept of horizontal and vertical time is also applicable to topical analysis and summarization. When applied to the entirety of the collection the result is a three-dimensional representation of the generated topics. Horizontal slices within this space represent a topical summary and segmentation of a slide, a lecture, or an entire course, based on the parameters requested. Vertical slices, through time, are also possible; generating a topical selection of relevant components through time.

Other dimensions can be added to a topic-based index of the eClass archive: author, discipline and course level are all possibilities.

5.3. Storage

Currently, the eClass archive consists of almost 3000 captured lectures. We have functioned under a “disk is cheap” paradigm for the eClass prototype. Accepting that paradigm for now, we are faced with two interconnected issues: *what should be stored and in what format*. If the resulting hyper-documents were fixed, we could answer the first issue by storing everything. However, many of the eClass components are fluid, inherently or in practice.

For example, the captured URLs are references to documents, which are outside the bounds of the eClass system. Capturing the entire HTML document is not feasible due to legal as well as storage reasons. Since a link can change, or the contents of the connected document may change, it is necessary to define and implement a method to determine the validity of a link. A possible draconian method would be to remove the link if the referenced page has been modified or tests as invalid. This can be tempered by producing a *signature* of the referenced material and only eliminating the link if there is a significant deviation.

There has also been a request by faculty to fix, remove, or replace erroneous material in past lectures. Currently, eClass only supports this manually. There is no support for preserving the past material. Changes leave no history and may alter the relationship between the components of the hyper-document.

Each of the components of an eClass document is stored in a separate format. The audio and video

components are stored as Real Media files, the slides are stored as GIF files and the ink data is stored in a file of its own. The relationship between components that make up a hyper-document is maintained physically within a directory for each document and logically within the structure of the database references. Voice transcriptions, teacher notes, and slide transcriptions are all stored within the database. In order to support a wider range of component sources, as well as higher quality media, the eClass system must be modified to support more media formats.

5.4. Other Challenges

The eClass system makes the archive available via dynamically generated web pages, as discussed previously. In its current implementation, eClass does not provide for access via media-limited devices like PDAs, cellular phones, or text only browsers commonly used for wearable computers. Like many other media rich archives, eClass hyper-documents lose significant content when reduced to one presentational data type. Though it is always an option to ignore the media-limited devices, eClass must contend with access requirements created by individuals who are perceptual impaired. Work needs to be done to make the visual elements of the archive available to the visually impaired, as well as making the audio components available to the audio impaired. The use of prosthetic technologies can assist, but cannot cover all of the media components represented in the eClass system [15].

Within the current subset of the population, who are supported by the existing interface, other accessibility issues arise. A significant issue is that of privacy. If those who were part of the original experience are the intended users of the captured archive, the issues of privacy are connected to the restriction of access to only those participants. However, if the captured experience becomes part of a more widely available archive, the problem of privacy becomes more complex. It may be argued that the archiving of a student’s likeness and voice, and the ideas or questions expressed by the student, are beyond the expectations of a student during the captured experience. Long-term use of the system is a growing trend. Students and faculty access captured experiences from previous semesters regardless of whether they were a part of that experience or not. The archive is accessed to acquire additional knowledge of a topic, to assist in creating future lectures, or to review past material. Since the archive is being used as a library of information, and not just as a set of electronic notes, individuals involved in the captured experience become collaborators in a multi-author document.

This also suggests the question of ownership within the archive. Capturing lectures and other experiences

raises questions of copyright and intellectual property rights. The question of ownership within the present system is largely ignored, by general consent. Faculty may request to have a course removed from the archive after the semester is completed. The vast majority, however, choose to leave their work within the archive, and allow it to be available for general access [12],[14].

6. Related Work

Capture of live experiences is an increasingly common theme in ubiquitous computing. Other research teams have used the notion of capture, integration, and access to facilitate collaborative or personal effort. Work at Xerox Corporation's Palo Alto Research Center (PARC) focused on capturing technical meetings to support summarization by a single scribe, who was often not well-versed in the subject of the meetings [17],[18].

Recently, the Informedia project at Carnegie Mellon University, initially set up as a digital library project to investigate sophisticated search capabilities on large bodies of video and audio, has been moving toward the ability to capture real-life experiences and search through them [31].

With the exception of the Tivoli work done at Xerox PARC[18] and Stifelman's Audio Notebook [25], very little attention has been paid to building a capture system for live use and evaluating its access use over an extended period of time. The Xerox PARC system observed the use of capture to support specialized meetings to discuss patent applications and the study lasted two years. Stifelman conducted two separate studies with different versions of her Audio Notebook and each study consisted of a small number of sessions both capturing and accessing personal notes from a variety of situations (class lectures, reporter interviews). Neither of these two projects has reported the scale of use reported here for eClass, mainly because their systems were not engineered to facilitate the automated collection and postproduction of materials as was the case for eClass.

Projects that are related to the educational application of eClass include the MANIC system at the University of Massachusetts [19], the DEPEND (Distance Education for People with Different Needs) system at the University of Oslo [21], the Chitra project at Virginia Tech [6], and the Lecture Browser system in the Project Zeno effort at Cornell University [32]. Other work includes the provision for more or less automated support to the capture of audio and video of the live sessions in order to have the presentation available for later review. This is the case of the ADEPT Project at Stanford [13] and the AutoAuditorium Project [10]. In such implementations, however, the information generated is mostly a linear video-based document.

The *Authoring on the Fly* system [9] provides an elaborate capture infrastructure that implements the linking by capturing paradigm. The implementation, which emphasizes the support to broadcasting the presentation live, supports capture of a fair number of applications running on Unix-only environments.

7. Conclusions

We have presented an overview of the eClass system, which provides the ability to automate the capture of live university lectures with a variety of short-term access interfaces. As the eClass repository grows, it presents a test bed for research in digital libraries. Some of these challenges include correct indexing and weighting of elements in the library, automated topic determination to facilitate search and summarization, appropriate storage decisions, universal accessibility and intellectual ownership and control of captured materials. With our existing educational repository, we will be able to explore many of these issues in detail. We expect that these results could be generalized to a variety of experience-based digital libraries.

8. Acknowledgements

The eClass project and its predecessor, Classroom 2000, resulted from the collaboration of many individuals. The authors would like to acknowledge the contributions of Maria da Graça Pimentel (University of Sao Paulo at Sao Carlos) and Yoshide Ishiguro (NEC Japan) for their specific contributions to the research that informed this paper during research visits with our lab. This work is sponsored in part by a grant from the U.S. National Science Foundation (CAREER grant IRI-9703384, Experimental Software Systems grant EIA-9806822, CISE Infrastructure grant EIA-9818305 and NSF/CNPq grant) and by grants from Sun Microsystems and Hewlett-Packard.

9. References

- [1] Abowd, G.D., Software Engineering Issues for Ubiquitous Computing, Proceedings of the International Conference on Software Engineering, Los Angeles, CA, May 16-22, 1999, pp. 75-84.
- [2] Abowd, G.D. Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment. *IBM Systems Journal*. Special issue on HCI / Pervasive computing, **38**(4):508-530, October 1999. . See <http://www.research.ibm.com/journal/sj/384/abowd.html>.
- [3] Abowd, G.D., C.G. Atkeson, A. Feinstein, C. Hmelo, R. Kooper, S. Long, N. Sawhney and M. Tani. Teaching and Learning as Multimedia Authoring: The Classroom 2000 project. *Proceedings of the ACM Multimedia'96 Conference*, Boston, MA, November 1996, pp. 187-198.
- [4] Abowd, G.D., C.G. Atkeson, J.A. Brotherton, T. Enqvist, P.A. Gully, & J. Lemon. Investigating the

- capture, integration and access problem of ubiquitous computing in an educational setting. *Proceedings of CHI'98*, ACM, pp. 440-447, 1998.
- [5] Abowd, G.D. and E.D. Mynatt. Charting Past, Present and Future Research in Ubiquitous Computing, *ACM Transactions on Computer-Human Interaction*, Special issue on HCI in the new Millenium, 7(1):29-58, March. 2000
 - [6] M. Abrams, S. Williams, G. Abdulla, S. Patel, R. Ribler, and E. A. Fox, Multimedia Traffic Analysis Using Chitra95, *Proceedings of ACM Conference on Multimedia*, San Francisco, CA (November 7-9, 1995). Also available as TR-95-05, Department of Computer Science, Virginia Tech (April 1995).
 - [7] Akscyn, R.M. and I.H. Witten. Report of First Summit on International Cooperation on Digital Libraries, June 1998, ks.com/idla-wp-oct98.
 - [8] D. Anderson, L. Harvel, M. Hayes, J. Jackson, and M. Pimentel. Internet Course Delivery – Making it Easier and More Effective, *Proceedings of International Conference on Multimedia and Expo*, 2000
 - [9] Bacher, C. and R. Muller. Generalized Replay of Multi-Streamed Authored Documents, *Proceedings of ED-Media '98*, Freiburg, June 1998.
 - [10] Bianchi, M.H., AutoAuditorium: A Fully Automatic, Multi-Camera System to Televis Auditorium Presentations. 1998 Joint DARPA/NIST Smart Spaces Technology Workshop, July 1998.
 - [11] Brotherton, J.A., G.D. Abowd and J. Bhalodhia. Automated capture, integration and visualization of multiple media streams. *Proceedings of the IEEE Multimedia and Computing Systems '98 Conference*, July 1998 pp. 54–63.
 - [12] Flanders, J., E. Mylonas. A Licensing Model for Scholarly Textbases, *Proceedings of the Fifth ACM Conference on Digital Libraries*, June 2-7, 2000, San Antonio, Texas, pp. 256-257.
 - [13] Harris, D., A. DiPaolo. Advancing Asynchronous Distance Education, *IEEE Transactions on Education*, August 1996
 - [14] Kumazawa, M., H. Kamada, A. Yamada, H. Hoshino, Y. Kambayashi, and M. Mohania. Relationship Among Copyright Holders for Use and Reuse of Digital Content, *Proceedings of the Fifth ACM Conference on Digital Libraries*, June 2-7, 2000, San Antonio, Texas, pp. 254-256.
 - [15] Luk, R., D. Yeung, Q. Lu, E. Leung, S.Y.Li and F. Leung. Digital Library Access for Chinese Visually Impaired. *Proceedings of the Fifth ACM Conference on Digital Libraries*, June 2-7, 2000, San Antonio, Texas, pp. 244-248.
 - [16] Mani, I. and E. Bloedorn. Summarizing Similarities and Differences Among Related Documents, *Information Retrieval*, Kluwer Academic Publishers, 1999, pp. 1-23.
 - [17] S. Minneman, S. Harrison, B. Janseen, G. Kurtenbach, T. Moran, I. Smith, and B. van Melle, "A Confederation of Tools for Capturing and Accessing Collaborative Activity," *Proceedings of the ACM Conference on Multimedia*, San Francisco, CA (November 7-9, 1995), pp. 523-533.
 - [18] T. P. Moran, L. Palen, S. Harrison, P. Chiu, D. Kimber, S. Minneman, W. van Melle, and P. Zelweger, "'I'll Get That Off the Audio,': A Case Study of Salvaging Multimedia Meeting Records," *Proceedings of ACM Conference on Human Factors in Computing Systems*, Atlanta, GA (March 22-27, 1997), pp. 202-209.
 - [19] J. Padhe and J. Kurose, An Empirical Study of Client Interactions with a Continuous-Media Courseware Server, Technical Report UM-CS-1997-056, University of Massachusetts (1997).
 - [20] Pimentel, M., G.D. Abowd and Y. Ishiguro. Linking by interacting: A paradigm for authoring hypertext. *Proceedings of Hypertext'2000*, May 2000, pp. 39-48.
 - [21] T. Plagemann and V. Goebel, "Experiences with the Electronic Classroom--QoS Issues in an Advanced Teaching and Research Facility," *Proceedings of the 5th IEEE Workshop on Future Trends of Distributed Computing Systems*, Tunis, Tunisia (October 29-31, 1997).
 - [22] Plaisant, C., B. Milash, A. Rose, S. Wildoff, and B. Shneiderman. Life lines: Visualizing personal histories. *CHI'96 Conference Proceedings*, pages 221-227, Vancouver, BC, Canada, 1996
 - [23] Richter, H.A., J.A. Brotherton and G.D. Abowd. Browsing Long-Term Captured Experiences, *Proceedings of CHI'00*, 2000
 - [24] Salton, G., A. Singhal, M. Mitra and C. Buckley. Automatic Text Structuring and Summarization, *Information Processing and Managemnet*, Elsevier Science, 1997, pp. 193-207.
 - [25] L. Stifelman, The Audio Notebook, Ph.D. thesis, MIT Media Laboratory, Cambridge, MA (1997).
 - [26] Swan, R. and J. Allan. Automatic Generation of Overview Timelines, *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2000, Athens, Greece, 2000.
 - [27] Takeshita, A., T. Inoue and K. Tanaka. Topic-based Multimedia Structuring. *Intelligent Multimedia Information Retrieval*, Maybury, M., ed., Cambridge, MA: AAAI/MIT Press.
 - [28] K. Weber and A. Poon, "Marquee: A Tool for Real-Time Video Logging," *Proceedings of ACM Conference on Human Factors in Computing Systems* (April 1994), pp. 58-64.
 - [29] Weiser, M., The Computer of the 21st Century, *Scientific American* 265, no. 3, 66-75, September 1991
 - [30] Whittaker, S. P. Hyland, and M. Wiley, "Filochat: Handwritten Notes Provide Access to Recorded Conversations," *Proceedings of ACM Conference on Human Factors in Computer Systems*, Boston, MA (April 24-28, 1994), pp. 271-277.
 - [31] See www.informedia.cs.cmu.edu/html/enter.html.
 - [32] See www2.cs.cornell.edu/zeno.