Developing an Educational Multimedia Digital Library:
Content Preparation, Indexing, and Usage

P. P. Semple and R. B. Allen
College of Library and Information Services
University of Maryland, College Park, MD 20742

A. Rose
Human Computer Interaction Laboratory and UMIACS
University of Maryland, College Park, MD 20742

ABSTRACT

The Maryland Electronic Learning Community (`www.learn.umd.edu`) is building a multimedia digital library of educational resources. Now in the fourth year of the project, we evaluate early decisions we made about segmenting and indexing videos. We also discuss an experiment in encouraging collaborative community indexing with a Quick Indexing Tool. We conclude that a broader base of users would better support the infrastructure requirements and we propose ways that such a broad base can be developed while also providing a framework for local learning communities. We propose a federated system of collaborative indexing communities.

1   INTRODUCTION

The Maryland Electronic Community (MELC) was designed to investigate how the professional development of teachers could be supported with information technology (Rose et al. 1999). One focus of this project has been the creation, maintenance, and cultivation of a digital library for the teachers to use as they develop interactive lesson plans (ILPs). While the original emphasis was on the needs of social studies and science teachers, in recent months the community has grown to include mathematics and language arts teachers.

The digital library contains over 3,200 multimedia resources, and by August 2000 this number is expected to exceed 6,300. The collection is comprised of text documents, audio files, video segments, still images, and Web sites. There are public domain items from the National Archives and NASA; proprietary items that have been contributed by project partners Discovery Communications and Maryland Public Television; and personal contributions from MELC members themselves. Each resource is identified by a set of standard metadata elements and then indexed according to Topic, Subtopic, and Standard (corresponding to national standards). Efforts at collecting and cataloging these resources was initially the responsibility of the team at the University of Maryland, but this work is gradually being handed off to the teachers as the external project funding ends.

By analyzing the Web logs, we can determine that the following resources were accessed for the first two months of the Fall 1999 school term: 53 images, 81 video clips downloaded, 9 audio clips, and 1 text file. Unfortunately, our server Web logs cannot show how many external Web sites have been accessed or how many videos were played via CD-ROM (videos can either be downloaded or played locally on CD-ROMs).

We are particularly interested in how the video clips are being used. Currently there are 52 videos in the library: 25 science and 27 social studies. Seven of these are in the public domain (from the National Archives and accessible by anyone) and 45 are private (from Discovery Communications, Inc. or

Maryland Public Television and accessible only by MELC members). Of the video clips downloaded, only 10 were public domain. The average video clip is about two to three minutes in length and approximately 15 MB. The Web logs show evidence of teachers downloading new videos from the server even before CD-ROMs are available. Putting the clips on the server reduces the time it takes to make the videos available to the teachers since creating and distributing the CD-ROMs is time consuming.

In this paper, we examine issues related to the use and development of the resource collection including:
- the flexibility of digitized video segments (as opposed to full-length film on the VCR),
- appropriate indexing and metadata elements, and
- collaborative indexing and the Quick Indexing Tool (QIT).

2   Segmenting the Videos

While the library contains images, audio, and video, we focus here particularly on the video because it is attractive to students, adds a unique dimension to lesson plans, and is invaluable to visual learners. The full-length educational videos are divided in one- or two-minute segments for a number of reasons. First, we consider the video segments to be building blocks for the teachers as they construct their lessons. Second, our technology was inadequate for the task of managing digitized full-length videos, especially at beginning of the project. A third reason for using segments is that they encourage cross-curriculum tie-ins. For example, "The Real Ben Franklin" is a film about the life of a great American and as a whole is clearly within the purview of the social studies department. However, it includes several fine illustrations of scientific principles and inventions that would be useful in a science class. As a result of the segmenting, science teachers can now use resources that they otherwise might not have been aware of or able to access.

While there has been a great deal of research in computer science on segmenting videos based on scene changes, we felt that truly meaningful units could be obtained only by manual segmentation. Unfortunately, while this kind of segmenting yields useful material, it is very labor intensive, requiring approximately seven minutes of processing for every one minute of video. In order to standardize the process, we have developed a set of guidelines.

As we have teased out the threads of these educational videos (interviews with experts in the field, historical footage, animations, demonstrations, re-enactments, first-person narratives, etc.), we have come to realize many ways that these segments might be re-constructed or regrouped by the teachers. Perhaps they will string together the re-enactments of Singbe Pieh and his fellow Africans aboard the Amistad, or pull out the animated graphics that show how black holes warp time and space. Thus, we try to segment along the different interviews or footage or speakers.

There were times when we chose not to segment the most discrete pieces of information, as when the sum of the whole was greater than the sum of the parts. This was most often when the medium itself, video, was providing something that could not have been captured and conveyed by other media. For example, near the end of the video "The Real Thomas Jefferson," the narrator stands on a balcony in Jefferson's American Room in Monticello. The camera pans around the room and as it passes over the artifacts the narrator briefly identifies them. Important though the images and the identification of these ancient Indian artifacts are, of greater importance is the panoramic view of the room itself. Similarly, we tended to give priority to the video's visual presentation of information over the text or narration that accompanied whatever was happening on the screen. Words can be gleaned from books or audiotapes, moving images come only from video.

We have developed several search tools that allow teachers to view these relatively small and discrete pieces. The most ambitious tool, the Explorer (Rose et al., 1995), gives teachers a visual representation of the entire collection, allowing them to select the data elements that will be displayed on the axes. With this tool, they can isolate all of the photographs from the African Americans in World War II photograph collection from the National Archives, or the resources that deal with aquatic habitats and that also satisfy the Personal and Social Perspectives objective of the National Science Education Standards (1995).

3   Metadata and Indexing by Content

Effective indexing is essential for any digital library and, in general, metadata and indexing should reflect the likely information needs of the library's users. Because this project was started several years ago and was focused upon the needs of a single school system, we had to develop much of our own metadata. The original indexing system was created by library-science students who tailored it to the needs of a specific group of middle school teachers. They reviewed science and social studies curriculum objects and textbook headings to determine what the list of topics and subtopics would be and compiled a lengthy set of instructions for future indexers. The result was a 60-page manual that outlined a system for assigning metadata elements (see www.learn.umd.edu/blcmanual.html) and specific details for assigning the descriptive elements: Description, Topic, Subtopic, and Standard.

Although a 60-minute video in its entirety may clearly belong in the social studies category, many of the segments from it could be cross-indexed to other topics. Unfortunately, many times these topics are far beyond the scope of our index. While it is probably a hallmark of high-quality educational resources that they use the familiar to introduce students to the unfamiliar, for the MELC community this is at once a blessing and a curse. As indexers, it leaves us with lots of segments for which we have no appropriate Topic or Subtopic, or which we can only index partially. The Discovery video entitled "Unearthing South America" is full of curriculum-related segments about early Indian civilizations, but many of these segments also deal with the culture's traditions of mummification and the archeological digs that are revealing much about these ancient people. Archeology is not a unit of study in the middle school curriculum, and so we have no way of capturing this essential data. It most cases we have adopted a "best fit" mentality and tried to shoehorn these segments into related categories.

A second stress upon the indexing system has been our growth as a community. As our ranks have swelled in recent years, we have embraced teachers from other areas, most notably mathematics and language arts. We have a handful of alternatives to accommodate this expansion. One would be to slow down the processing of videos and take the time to construct an index that is similar in scope and detail to the indexes for science and social studies. However, not only is this a very labor-intensive task, but it would suggest that we have learned none of the lessons of our own recent history (what if the foreign language teachers join us next month?). Our discussions centered on two other viable options: (1) distill the descriptive indexing down to the most basic and community-specific of categories and simply assign segments to the curriculum unit(s) where they are most likely to be used; or (2) broaden the descriptive indexing structure to encompass all possible subject areas by using a scheme such as Dewey, which in theory covers the entire body of knowledge.

In retrospect, we might have limited our index to the broadest of categories for the subject areas, including many beyond the curriculum, and emphasized techniques for keyword searching in the Description field. Concurrently, we would have enriched the descriptions. Rather than simply giving a brief one- or two-sentence summary, we would have made a greater effort to get the transcripts for the videos and simply copied the relevant passages directly into the Description fields, thus providing the teachers with richer text to mine. Regardless, we have come to discover that having the transcripts for

any sort of video indexing is highly desirable. Short of taking the time to implement a controlled vocabulary or develop an authority file, it is important to refer to the same entity in a consistent fashion. For example, a recent addition to the library has been a video on HIV and AIDS. In a system based on simple character string recognition, it is imperative that the teachers know if the film presents information on "CD4 T cells," "cd4 t cells," "C D 4 T-Cells." or "see, de four tea sells"?

Our custom-built indexing effort may not be sustainable. We feel our efforts would be better spent indexing to some national metadata standard such as the Dublin Core or GEM (Sutton, 1999). While we do index to the National Science Education Standards and the Curriculum Standards for Social Studies, we do not have sufficient data to suggest that this field is being used by the community. Similarly, we have considered searching for an indexing system that would apply beyond our middle school teachers to a high school or elementary school community.

4   Collaborative Indexing

Our goal has been for MELC to become a self-sustaining community when the project funding ends. At the heart of this community of educators would be the digital library that supports their work in the classroom. In order for this project to be self-sustaining within the current framework of the Baltimore City Public School System, we believe that the indexing of new resources would have to be done collaboratively by members of the community.

While our effort is on a much smaller scale than the CORC project (OCLC, 1999), we certainly share their vision. Perhaps there are even some parallels between the indexing of the multimedia content and the manner in which Varian (1999) has suggested that quality control in scholarly publishing can become self-regulating.

4.1   Quick Indexing Tool (QIT) and Collaboration

Our first effort at fostering a system of collaborative indexing, the Quick Indexing Tool (QIT) (209.48.188.166/blclib/quickindex/), is designed to let the teachers capture descriptive and bibliographic information (metadata) about Web sites as they add them to the digital library. Having come to know both the capabilities and needs of our users, we have designed a simple interface that lets them capture the greatest amount of information with the least amount of effort.

Along with several simple fields for each QIT submission, we have included a description field that encourages teachers to jot down a brief statement or vocabulary from the page itself (keywords) as a way of representing the content. In the second half of the form, we ask the grade, subject area (social studies, math), and unit of study (Baltimore uses the Berri Curriculum) for which they envision this resource will be useful. We anticipate that these particular required fields will not dampen enthusiasm for the tool by implying that all teachers should be expected to know about the specifics of all units of study across all grade levels and all subjects. We anticipate that sixth-grade math teachers will be entering Web sites applicable to the sixth-grade math curriculum and seventh-grade science teachers will be entering Web sites applicable to the seventh-grade science curriculum units, and so forth. Having the resources selected by these subject experts should greatly increase the usefulness and quality of the library.

The QIT was introduced to the MELC teachers in August 1999. To date, 28 resources have been added using this tool. It has helped us expand the library to include the new content areas of math and language arts by having the new teachers themselves make contributions. Each of the 11 math resources in the library have been added by math teachers who used the QIT.

4.2. Pros and Cons of Collaborative Indexing

In a system of collaborative indexing, all community members would be responsible for resource acquisition and cataloguing. Resources can be found items, such as Web sites; proprietary items, such as the Discovery video segments; or self-generated items, such as the photographs of Baltimore and Europe taken by various teachers who scanned them and then entered them as digital images. The benefits of such a system are obvious and important:

1) In theory, it guarantees that each item will be have high relevance to some subset of users. Only resources that other members of the group have identified as being specifically related to the curriculum will be part of the collection.

2) Not only will the items all have high levels of relevance, but they will also be of high quality. Each item will have been judged useful by a subject expert.

3) Moreover, the activity of creating a QIT entry may lead to reflection by the teacher about the value of the item. The time involved in entering the bibliographic data, which includes their name and e-mail, will serve as further impetus for maintaining quality control.

However, there are a few difficulties in expecting teachers to use the QIT on a regular basis:

1) Becoming a contributing member of a digital library community means more work for a group that is habitually overworked. Society, school administrations, and parents already expect them to perform a myriad of functions within their classrooms. We are asking them to perform yet one more.

2) The community of users probably will not find as many resources as a library media specialist would. Although the community users will be able to index the resources they do find, they may be overlooking many others. For example, the media specialist knows that many valuable sites are only available to users who know to execute a search on the main page or by searching the site index, and that these pages cannot be found through the major search engines. Unlike today's librarians, teachers have not necessarily been trained to manipulate the search engine algorithms as a way to refine the search process.

3) The community users probably will not find resources of the same quality as a media specialist. The reason for this is similar to the argument made in the previous point.

4) The community of users and the QIT will not exist in an equilibrium. As we have discovered, changes in the environment can have messy implications for the system as it was originally conceived. If they are not provided with the technical skills or training to respond to these changes, then the slightest shift in the equilibrium may well overwhelm them.

5   Conclusion

Based on our experience, we envision an extended collaborative indexing system that is overseen by a professional librarian with whom they have ready access and easy communication. This professional librarian would serve as a resource to be used by the community members on an "as needed" basis, although this professional will also be an equal contributor to the endeavor and a person who will no doubt have many useful things to share when it comes to the selection, acquisition, organization, and use of resources. We envision an iterative process:

- an individual or subset of members approaches the media specialist with a request for information,
- the media specialist will interview them to further understand the nature of their information need,
- the media specialist will provide the user with a list of resources,

- the teacher will examine the resources,
- the teacher will select and index (with QIT) the ones that meet his/her need, including those that he/she may have found on their own.

This scenario strikes a balance between two extremes. It leaves ultimate responsibility for resource selection to the community members, thus ensuring high relevance and importance to the curriculum, but it lets them have the option of having input from a professional whose intervention represents a great deal of "value added" to the final product.

Finally, it also suggests a way that we can provide a useful framework for the local learning community, yet provide an infrastructure for a broader community. Once a media specialist has compiled a listing of resources that would be useful for an eighth-grade lesson plan on meteorology, it is a relatively simple task to share that list with other teachers in other local learning communities.

Perhaps the solution lies in a federated system of collaborative indexing communities. In this arrangement, many small groups of collaborative indexing communities join forces to benefit from the sharing of resources, expertise, and support that are inherent in a larger federation. Suppose an individual member of Community A has annotated and cataloged a resource in a way that has made it useful to the specific needs, parochial though they be, of the members of her immediate community. An individual in Community B, upon perusing the digital library holdings of Community A, accesses this same resource but represents the descriptive data differently in his catalog. For his users, he includes information about how the resource correlates to statewide standards outlined by the governor's task force. On a macro level, the members of Community B benefit from having access to the resources in the digital library of Community A; on a micro level, they benefit from the refining process that defines collaborative indexing systems.

We have plans for adding tools similar to QIT that will allow teachers to index video segments, materials from online archives, and other multimedia items. This will be accomplished by redefining the metadata fields in the top half of the form so as to prompt users to give more resource-specific information (e.g., audio tapes often have narrators, computer images come in varying file formats).

We continue to believe that educational video can be highly absorbing and can be an important resource for middle school teachers. However, many of the costs for the digital library (indexing, segmenting) would be better borne by a very large constituency. The issue becomes maintaining a sense of involvement from the local learning communities. The QIT is one example of a community-oriented tool. Other such tools might include social collaboration tools (see Allen, 1997).

6 Acknowledgements

7 References

Allen, R. B., "Mental models and user models." *Handbook of Human-Computer Interaction* (2nd edition) Edited by: M. Helander, T.K. Landauer, and P. Prabhu, Elsevier Science, 1997, 49-63.

OCLC, "Cooperative Online Resource Catalog (CORC)", 1999, http://www.oclc.org/oclc/research/projects/corc/index.htm

MPEG7-7 Context and Objectives, http://www.darmstadt.gmd.de/mobile/MPEG7/.

Rose, A., Ding, W. Marchionini, G., Beale, J., and Nolet, V. "Building an Electronic Learning Community: From Design to Implementation." *Proceedings ACM CHI,* 1998, Los Angeles, 203-210.

Rose, A., Allen, R. B., and Fulton, K. "Multiple Channels of Electronic Communication for Building a Distributed Learning Community." *Proceedings Computer Supported Cooperative Learning*, 1999, Palo Alto, 495-502.

Sutton, S. A. "Conceptual Design and Deployment of a Metadata Framework for Educational Resources in the Internet." *Journal of the American Society of Information Science*, 1999, 1182-1192.

Varian, H. R., The Future of Electronic Journals, **Error! Bookmark not defined.**