Int J Digit Libr (2004) 4: 36–38 / Digital Object Identifier (DOI) 10.1007/s00799-003-0058-4

A briefing on the evolution and status of the Open Video digital library

Gary Marchionini

100 Manning Hall, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA
e-mail: march@ils.unc.edu

Published online: 27 July 2004 –

Abstract. This briefing introduces the reader to an ongoing digital library project that serves as a testbed for research that in turn informs the development of the digital library. The paper is informal with pointers to the technical literature. It aims to illustrate some of the issues and challenges facing digital library researchers and practitioners.

Keywords: Video retrieval – Interactive interfaces

Digital libraries have evolved out of traditional library technical services and the recognition by the computer science community of the importance of new applications of data models that integrate DBMS and IR traditions. Libraries have been automating catalogs, circulation, and other services since the 1960s. Because the library marketplace is miniscule compared to banking, manufacturing, telecommunications, and other industries, library automation progress was dependent on developments in those industries throughout the 1970s and 1980s. Large investments in scientific databases, the development of the Internet, and mass-market computing caused some research policy leaders to suggest that large-scale citizen access to data could be a grand challenge that would continue to spur technical development while returning benefits to society. Computer scientists and librarians rose to the challenge of applying their research and development interests to this challenge under the umbrella of digital libraries (DLs). With this as context, this author's work in digital libraries evolved from core interests in how people seek and use information in electronic environments to evaluating and building digital libraries. In 1990, my work in evaluating the nascent Perseus Project (www.perseus.tufts.edu) made it clear that what was developing was a DL [5, 6]. That work continued for a decade. Concurrent with these efforts my colleagues and I worked with the Library of Congress National Digital Library Program to determine user needs and create prototype user interfaces [9] and with the Baltimore City Public Schools to develop a video DL for middle school science and social studies classrooms called the Baltimore Learning Community (BLC) [8]. The BLC Project was termed a "sharium" in 1997 because it encouraged teachers to share their lesson plans and experiences in using these digital materials, i.e., to become active participants in the digital library (see [4] an elaboration of the sharium concept). The experience of the BLC led to the definition of the Open Video Project, an open-source digital video repository and testbed for the research and educational communities [12].

The Open Video Digital Library (OVDL) aims to provide the research and educational communities with open-source digital video content while serving as a testbed for DL research and development (www.open-video.org). This briefing provides an overview of the project with attention to the interplay between our research goals and implementing a reliable production system. As of this writing, in the fall of 2003, the OVDL references more than 1800 video segments (more than half a terabyte) and is accessed by about 5000 unique visitors each month. Most of the video is available as MPEG-1 files, although MPEG-2, MPEG-4, and Quicktime files are available for some segments. An important principle in the sharium spirit is to provide files that are downloadable and manipulable rather than only video streams that are mainly limited to display. The files are distributed via the Open-Video Channel on the Internet2 Distributed Storage Initiative Network, and the metadata repository resides on a server at UNC-Chapel Hill. The repository consists of a MySQL database with 12 tables that include 84 fields (including all keys and foreign keys). The

scheme is a Dublin Core superset and is exposed to OAI harvesting. The primary user interface is PHP-driven and provides multiple entry points including: free text search (query form); controlled vocabulary search (attribute selections and pull-down pick lists); and browsing by genre, duration, color, sound, and contributing organization. Once a partition of the database (hit list) is identified through search or browse functions, users can select full bibliographic details, storyboard previews, fast forwards through the entire clip at 64X, or a 7-s extract or download the segment. People or organizations can contribute video to the OVDL in several ways. Some popular materials include the University of Maryland's Human-Computer Interaction Laboratory videos and the ACM CHI videos. See [7] for a recent paper describing the details of the OVDL.

A key challenge in the OVDL is to build and maintain a functioning production system useful to the public while conducting research in user interfaces for video retrieval. A functioning production system must be available 24/7 and handle updates seamlessly; in addition, it depends on a stable and well-defined workflow process. On the other hand, research environments must be highly flexible to allow edge-of-envelope experiments and prototype testing. Our approach to these opposing requirements is to build a production component and use it as the basis for a research component that allows us to plug in different UI modules or backend procedures and evaluate discrete elements in our usability lab. Over time, as we test new features, these features will be integrated into subsequent releases of the production system. For example, the first version of the production system provided storyboard overviews and previews for video segments as users clicked on tabs and/or links. Our UI research goal is to build and test agile views that give people multiple representations for information and intuitive control mechanisms that facilitate easy switching among the views [10]. In our experiments, we created slide shows with and without audio keywords, storyboards with and without audio keywords, and a variety of fast forward surrogates. These surrogate prototypes were created for a subset of the full database, and laboratory user studies were conducted to determine how they influence search and sense making. An exploratory comparison of surrogates [13] convinced us to drop the slide shows and develop fast forwards. A subsequent comparative study of different fast forward speeds [14] led us to select 64X as the fast forward rate rather than 32, 128, or 256 rates. Studies of shared and history views [1] led to the incorporation of recommendations and extensive mouse hover mechanisms to move between views. Another study used eye tracking to study the effects of placing poster frames in results lists (number of frames, integration with textual metadata) and demonstrated the importance of text titles for gist determination [2]. These studies are driven by an overall set of research questions related to the interaction of video and user characteristics, tasks, and surrogates. The idea is to discover useful surrogates and their boundary conditions (for setting user-selectable settings and system defaults) that can be integrated into future versions of the production system.

In addition to the main research interests in UIs and surrogation, we are also investigating other DL problems such as metadata management and user contributions. Our approach to metadata has evolved from a simple Dublin Core set of elements to a more robust set of optional elements in the 12 tables noted above. These elements map fairly well to other approaches such as the ongoing discussions within the Public Broadcasting Metadata Dictionary Project. We actively seek to acquire video from important sources and support user contributions to OVDL in several ways. People may use a form-based interface to describe a video and submit either or both of the metadata (including a stable address for access) and/or the primary video files. Submissions are added to the database automatically; however, they are not made publicly available until an OVDL curator approves release. The approval policy is at present rather simple as we have not encouraged or received random contributions. The more typical approach to contribution is for people or organizations to provide tapes or files and work with us to create metadata records. The digital files can reside on the contributor, on the OVDL servers, or on both. Key examples of video contributed in this collaborative manner include documentaries from the Intermedia Project at Carnegie Mellon University, ephemeral films from the Prelinger Archive, symposium videos from the University of Maryland Human-Computer Interaction Laboratory, ACM CHI videos, and documentaries from NASA. Finally, we also crawl publicly accessible reliable sites (e.g., Library of Congress) and harvest the metadata and add them to the OVDL automatically.

The underlying asset management process depends heavily on careful attention to the underlying MySQL database. For example, although new content can sometimes be piped automatically, attention to broken links and adding new surrogates for the files are currently ad hoc processes. As we gain more experience, the backend processes of managing a video DL become more routinized. We have developed a variety of tools and scripts in this regard. We have modified the MERIT system [3] to extract keyframes on our Linux platform, created programs to create fast forward surrogates for each clip (according to a naming and directory structure scheme), experimented with procedures to turn text keywords into speech-synthesized keywords, and developed methods for digitizing analog tapes and saving resulting uncompressed files in various file formats. Additionally, we have developed a metadata browser tool (VIVO), a peer-to-peer tool to facilitate sharing of video files among digital librarians, and a tool for using digital video in distance education settings (ISEE, see [11]). Ultimately, we aim to integrate the various scripts, tools, and procedures into an open-source digital video toolkit.

Our goals are to continue to develop new surrogates and add substantial content. These upgrades will incorporate more visual surrogates (e.g., dynamic poster frames, more audio cues, and shared views based on collaborative filtering analyses). Research on surrogates that automatically incorporate more speech into text and feature detection advances is also a high priority.

In sum, the OVDL represents ongoing efforts to build and test an open-source digital video digital library that can meet the needs of a variety of audiences and serve as a testbed for digital video retrieval and DL research. The OVDL is meeting both of these goals on multiple fronts. The TREC video track uses OVDL content as part of its testbed for video retrieval, and a number of research groups around the world use the content in their research. We participated in the 2003 TREC video interactive search track. Several educators are using the content in classes to illustrate points or as part of larger presentations, and students also incorporate clips into their work. The repository has served our UI research and development well by providing a variety of genres and a critical mass of content for creating indexes, surrogates, and interfaces to test our design framework. Thus the OVDL is both a working DL accessed by people around the world for research and educational purposes as well as the seed of the larger sharium notion where people come virtually to get and contribute resources. As it evolves, the tools as well as resources will hopefully also make it a conceptual place where people collaborate and work.

Acknowledgements. This work is supported by NSF Grant #IIS 0099638.

References

1. Geisler G (2002) AgileViews: a framework for creating more effective information seeking interfaces. Unpublished PhD dissertation, University of North Carolina at Chapel Hill
2. Hughes A, Wilkens T, Wildemuth B, Marchionini G (2003) Text or pictures? An eyetracking study of how people view digital video surrogates. In: Proceedings of the 2nd international conference on video retrieval (CIVR 2003), University of Illinois, Champaign-Urbana, 24–25 July 2003. Lecture notes in computer science, vol 2728. Springer, Berlin Heidelberg New York, pp 271–280
3. Kobla V, Doermann D, Faloutsos C (1998) Developing high-level representations of video clips using VideoTrails. In: Proceedings of the SPIE conference on storage and retrieval for image and video databases VI, San Jose, CA, 24–30 January 1998, pp 81–92
4. Marchionini G (1999) Augmenting library services: toward the sharium. In: Proceedings of the international symposium on digital libraries, Tsukuba, Japan, 28–29 September 1999, pp 40–47
5. Marchionini G (2001) Evaluating digital libraries: a longitudinal and multifaceted view. Libr Trends 49(2):304–333
6. Marchionini G, Crane G (1994) Evaluating hypermedia and learning: methods and results from the Perseus project. ACM Trans Inf Sys 12(1):5–34
7. Marchionini G, Geisler G (2002) The open video digital library. dLib Mag 8(12). http://www.dlib.org/dlib/december02/marchionini/12marchionini.html
8. Marchionini G, Nolet V, Williams H, Ding W, Beale J, Rose A, Gordon A, Enomoto E, Harbinson L (1997) Content + Connectivity ⇒Community: digital resources for a learning community. In: Proceedings of ACM Digital Libraries '97, Philadelphia, 23–26 July 1997, pp 212–220
9. Marchionini G, Plaisant C, Komlodi A (1998) Interfaces and tools for the Library of Congress National Digital Library Program. Inf Process Manage 34(5):535–555
10. Marchionini G, Geisler G, Brunk B (2000) Agileviews: a human-centered framework for interfaces to information spaces. In: Proceedings of the annual meeting of the American Society for Information Science, Chicago, 12–16 November 2000, pp 271–280
11. Mu X, Marchionini G (2002) Interactive shared educational environment (ISEE): design, architecture, and user interface, April 2002. SILS Technical Report: TR-2002–09, UNC-Chapel Hill. http://ils.unc.edu/idl/isee/ISEE_tech_report.pdf
12. Slaughter L, Marchionini G, Geisler G (2000) Open Video: a framework for a test collection. J Netw Comput Appl (Special Issue On Network-Based Storage Services), pp 219–245
13. Wildemuth B, Marchionini G, Wilkens T, Yang M, Geisler G, Fowler B, Hughes A, Mu X (2002) Alternative surrogates for video objects in a digital library: users' perspectives on their relative usability. In: Proceedings of the 6th European conference ECDL 2002 on research and advanced technology for digital libraries, Rome, 16–18 September 2002. Springer, Berlin Heidelberg New York, pp 493–507
14. Wildemuth B, Marchionini G, Yang M, Geisler G, Wilkens T, Hughes A, Gruss R (2003) How fast is too fast? Evaluating fast forward surrogates for digital video. In: Proceedings of the ACM/IEEE joint conference on research on digital libraries Houston, TX, 27–31 May 2003. IEEE Press, Los Alamitos, CA pp 221–230