Interfaces and Tools for the Library of Congress National Digital Library Program

Gary Marchionini, Catherine Plaisant, and Anita Komlodi

University of Maryland Human Computer Interaction Laboratory and Digital Library Research Group

{march, komlodi}@glue.umd.edu, plaisant@cs.umd.edu

Abstract

This paper describes a collaborative effort to explore user needs in a digital library, develop interface prototypes for a digital library, and suggest and prototype tools for digital librarians and users at the Library of Congress (LC). Interfaces were guided by an assessment of user needs and aimed to maximize interaction with primary resources and support both browsing and analytical search strategies. Tools to aid users and librarians in overviewing collections, previewing objects, and gathering results were created and serve as the beginnings of a digital librarian toolkit. The design process and results are described and suggestions for future work are offered.

Digital Libraries (DL) offer new challenges to an emerging breed of digital librarians who must combine the principles and practices of information management with rapidly evolving technological developments to create new information products and services.

This paper describes a collaborative effort to explore user needs in a digital library, develop interface prototypes for a digital library, and suggest and prototype tools for digital librarians and users at the Library of Congress (LC). Over a two-year period, a team of designers from the University of Maryland worked with LC staff to develop the interfaces and tools for the LC National Digital Library Program. The interfaces were designed according to the principles that users should maximize their interactions with information resources and minimize their attention to the system itself, and that both browsing and search strategies should be supported. These principles were manifested in the following specific design goals:

- Minimize disorientation by reducing navigation (e.g., minimize scrolling and jumping, flattening hierarchy) and anchoring users in a consistent context;
- Provide primary information at the earliest point in the interaction as possible;
- Support rapid relevance decisions through overviews and previews.

The tools were designed to help users and LC staff to examine and use different digital objects. There were three types of tools developed:

- Overviews of collections
- Object "previwers"
- Object "gatherers"

This paper provides background on digital libraries and the LC National Digital Library Program, describes the overall collaborative project, focuses on the iterative design of the interfaces and tools for a special collection, and concludes with a discussion of research and development needs.

Digital Library Research and Development

Digital libraries are the logical extensions and augmentations of physical libraries in the electronic information society. As extensions, digital libraries amplify existing resources and services and, as augmentations, they enable new kinds of human problem solving and expression. High levels of attention and funding were first given to digital libraries in the early and mid 1990's, leading to a plethora of visions and projects invariably driven first by finding ways to apply the many technologies developed in the 1980's and second by desires to create new technologies for managing distributed information resources (see April 1995 issue of Communications of the ACM for descriptions of various DL projects and the May 1996 issue of IEEE Computer for descriptions of six prominent projects). In some cases, notably the efforts of national libraries or large academic libraries, efforts focused on extending access to existing collections through digitization and network access. At present, digital library projects have focused on developing digital collections and providing rather limited access services.

Marchionini (in press) has characterized DL research and development as falling into four categories: content, services, technology, and culture. Research issues related to content includes the integration of multimedia objects; data acquisition, including analog to digital conversion; metadata extraction and standardization; procedures for indexing, storage and retrieval; workflow processes and management; and collection preservation and maintenance. Service research issues are strongly dependent on user interfaces and include search, filtering and browsing; reference and question answering; and instruction.

Technology research efforts are mainly related to high-speed networking, search engines, interface design, security and billing, and interoperability across many DLs. The culture issues include intellectual property; insuring data quality, privacy, and equity; and organizational interfaces for various communities of practice. In addition to these research and development challenges, meta issues related to managing and evaluating DLs and their impact on people and organizations are also active areas of study.

The Library of Congress (LC) has long served as an exemplar for library practice and led many institutions in defining library automation needs. In 1995, LC initiated an ambitious National Digital Library Program that aims to digitize and make available five million objects from a variety of LC collections. Funded by corporate and foundation donations, the program will not only develop one of the premier digital libraries for American culture, but also provide an opportunity for researchers and practitioners to address many technical and intellectual challenges that DLs provide. The program is pushing the state of the art in digitizing large volumes of non-book and sometimes fragile objects, creating high quality packages for K-12 classrooms, breaking new ground in corporate-government partnerships, operationalizing sophisticated search services, and creating innovative interfaces for public and staff users (see www.loc.gov). In 1995, the LC NDL Program contracted with the University of Maryland's Human-Computer Interaction Laboratory to collaborate on the interface development for the NDL.

The focus here is on user interfaces and the work of researchers creating interfaces for information retrieval and data visualization is particularly applicable to DLs. Although there have been a number of innovative approaches to supporting search, the WWW environment has only recently begun to support the basic interaction techniques available in standalone or earlier client-server environments (e.g., multiple windows, mouse dynamics, stateful interaction, etc.). The extension of HTML (e.g., frames) and especially the development of Java are allowing designers to port interfaces to the web environment so important to DLs. For example, several projects aim to improve user information seeking by closely coupling text search with visualization methods to display the result of the searches (e.g. Tilebars [Hearst, 1995], Envision [Heath et. al., 1995]). Other projects allow users to directly manipulate data, thus integrating search and display, (e.g., dynamic queries [Ahlberg & Shneiderman, 1994], continuous zooming [Bederson & Hollan,], overlays and lens' [Rao, et.al., 1995].

DL and multimedia research has led to innovative interface widgets for specialized materials (e.g., video skims [Wachler et. al., 1996], query previews [Doan et. al., 1996], and zoomable maps [Smith, 1996]). Improved tools have spurred human-computer interaction researchers to apply interface design metaphors and widgets to web browsing applications (e.g., book metaphor [Card et al., 1996], decks of cards metaphor [Brown and Shillner 1996], tiled, elastic windows [Kandogan and Shneiderman 1996]. Thus, today's DLs that depend on WWW are increasingly able to take advantage of the fruits of human-computer interaction research.

In our NDL project, there were two basic decisions we made to apply dynamic query interaction styles to meet the project goals. First, at the project's inception we adopted

JavaScript and Java to overcome the limitations of HTML display and cgi script-based interaction. Second, to overcome the lack of common metadata for all objects in the NDL, we divided user browsing into two components by developing a collection browser distinct from the within-collection browser. These solutions are illustrated in the resulting designs.

Phase 1: Problem Identification and Team Development

The first phase of the project entailed defining the NDL interface context, developing a user-centered interlibrary design team, and identifying the basic functionality for NDL interfaces. In this phase (September 1995 through May 1996) the following activities took place:

- A NDL working group was formed, a team mail reflector and web site were established, and regular meetings were held to brainstorm, sketch designs, and react to prototypes;
- A series of four public briefings were held at LC to focus attention on key user-centered design topics;
- A user needs assessment was conducted that elicited input from LC Reading Room staff, teachers involved in NDL orientation sessions, school library media specialist supervisors in the state of Maryland, and parents and workers at a day-care facility in Michigan;
- A series of interface designs were sketched and mocked up for discussion.

Basing interface designs on the needs of users is a central principle of user-centered design. A number of steps were taken to assess the needs of traditional users of LC collections and to project the needs of potential users attracted by the remotely accessible NDL. The user needs assessment included interviews with LC librarians in nine Reading Rooms at LC, a questionnaire sent to teachers from school districts around the United States who participated in developing and testing LC's Learning page, a questionnaire sent to all school library media specialist supervisors in the state of Maryland, a questionnaire distributed to parents and employees at a large day-care facility in Michigan, and document analysis of various reports and handouts produced by LC staff (see Marchionini, Plaisant & Komlodi, 1996 for details on this phase of the project). The results illustrated the need for content-specific search strategies as well as pose a set of design challenges across all the collections:

- Serving a wide range of users
- Serving a wide variety of information needs
- Helping users distinguish primary and secondary materials (including multiple layers of each)
- Helping users make links among items across different collections and reading rooms
- Capturing the essential elements of the reference interview so that users can find what they need without human intervention
- Communicating to the user what items are NOT in the NDL and giving pointers to external resources
- Creating an interface that is accessible by users with state-of-the-market technology
- Developing new techniques to search for multimedia objects and to integrate those

techniques into the interface (e.g., visual and audio query languages)

- Providing clear delineations within the NDL as well as among the larger collections of the physical Library of Congress and the entire Internet
- Helping users distinguish secondary and primary materials
- Distinguishing and rapidly displaying texts, graphics, and sounds to facilitate browsing
- Allowing users to search across collections or not; and if they choose to limit searching to a collection, provide clear linkages to other collections

Meeting all these challenges will ultimately take many iterations of interface implementation and testing, but the aim of the project was to address as many as possible in the second phase that entailed prototyping interface designs and digital librarian tools.

Phase 2: Interface Design

The second phase (June 1996-August 1997) used the experience of the first phase to focus the project on two goals. First, the project team aimed to apply the mockups from phase one to an integrated prototype that illustrates solutions to the high-level design goals and specific NDL challenges. Second, the project aimed to develop and integrate into the prototype specific overview and preview tools to support different media types and user activities. These tools could be used by end users or by librarians building the NDL. To better model the interface prototyping and tool development, we focused on a specific collection (Coolidge Consumerism) that draws materials from multiple LC collections and includes multiple media types. The integrated interface prototype gave users entry into the entire NDL (termed the American Memory level--AM) and illustrated a coordinated, within-collection interface for the Coolidge Consumerism (CC) corpus.

Interface Prototypes.

One of our goals was to *reduce navigation* to save time and also to reduce the possibility of disorientation. Disorientation was a well-known side effect of hypertext navigation before the WWW emerged (e.g., Conklin, 1987, Nielsen, 1990) and it is clearly an issue in today's web environment. Disorientation is mainly due to the cognitive load necessary to navigate a conceptual space with few unique landmarks. A related side effect is distraction as users follow seductive links rather than the logical progression of a theme that spans multiple hypertext nodes. Although many researchers are working on ways to physicalize abstract information spaces with visualization techniques such as 3-D views and virtual reality walkthroughs, a more practical approach was taken in this project. The design philosophy adopted was to reduce navigational requirements by flattening the hierarchy and anchoring the user in a stable context.

Broad and shallow menu structures have been shown to be more effective than narrow and deep structure (e.g., Landauer & Nachbar, 1985). Our designs attempted to *flatten the hierarchy* of the site by bringing to the surface many samples of items and special presentations, using frames and overviews such as tables of contents to reduce the need for backtracking in the page structure and by offering basic search facility at the top of the hierarchy.

Our prototypes also illustrate the benefits of *anchoring users in a stable context* with a consistent design. The interfaces for both the American Memory levels and the individual Coolidge collection are coordinated through the used of a similar layout and organization. A standard toolbar on the left of the screens lists the functions available. A goal was to illustrate that users could get a good sense of what the collections contain (samples), browse and select collections and finally search and review results without leaving the home page. This approach was meant to minimize user disorientation and distraction and allow users to focus their cognitive efforts on the NDL resources rather than system navigation.

The design also emphasized providing *previews and overviews* of materials to users (Greene et. al., 1997). Previews and overviews have several functions in digital libraries. They help users learn about the collection and explore the materials along different attributes which in turn can help support searching by enabling users to better define, limit and focus their searches. The American Memory digital collections contain very large amounts of materials in many different formats with varying levels of descriptive metadata. This makes searching difficult and browsing more important. Previews and overviews support browsing by allowing users to quickly make judgments about the efficacy of examining full versions of objects and in the affirmative case providing rapid access to those objects.

Previews are extracted from the primary data objects, they represent specific collections or items. Overviews are constructed from large sets of items and show different attributes of these items and aggregate information about the set. A preview of one collection can be samples from that collection, while a preview of a photograph can be a thumbnail. An overview of all the collections shows on one screen each of the collections with some of their attributes. An overview of all the items in a collection shows on one screen a representation of all items, possibly with some attributes. Looking at previews of result sets can speed up the evaluation process and save users time. Both previews and overviews help users learn about the collections and get a better sense of what is available and just as importantly, what is not available (for example by noting that many primary objects are not yet digitized).

The application of these principles is illustrated below. The examples are organized into layout consistency across the AM and CC coordinated interfaces and by main functions of these interfaces (introduction and samples, browse, and search).

Consistent Layout. The main functions of the interface are arranged in "tabs" down the left side of the screen on a toolbar that is always present. Selecting a “tab” on the left brings the corresponding details on the right (using HTML frames). This design (See Figure 1a) is more compact than our previous iterations and leaves space for the introduction and the samples to be visible when users first see the page. Both interfaces (American Memory [AM] and Coolidge Consumerism [CC]) have consistently placed tabs for “introduction” (selected by default), “Browse”, “Search” and “Help”. The tabs at the top reflect the LC hierarchy: from the main LC home page, to American Memory, down to the collection in the case of the CC interface. Although frames have been justly criticized as problematic for bookmarking, printing, and backtracking functionality, we believe that well-designed anchoring frames serve to help users maintain cognitive balance and that this advantage far outweighs the possible side effects of frames.

The two interfaces are differentiated by their color scheme (blue for AM and red for CC, see Figure 1b). Of course a single color is not sufficient for all the collections, but the current collection home pages already demonstrate that different styles can be devised for the different collections.

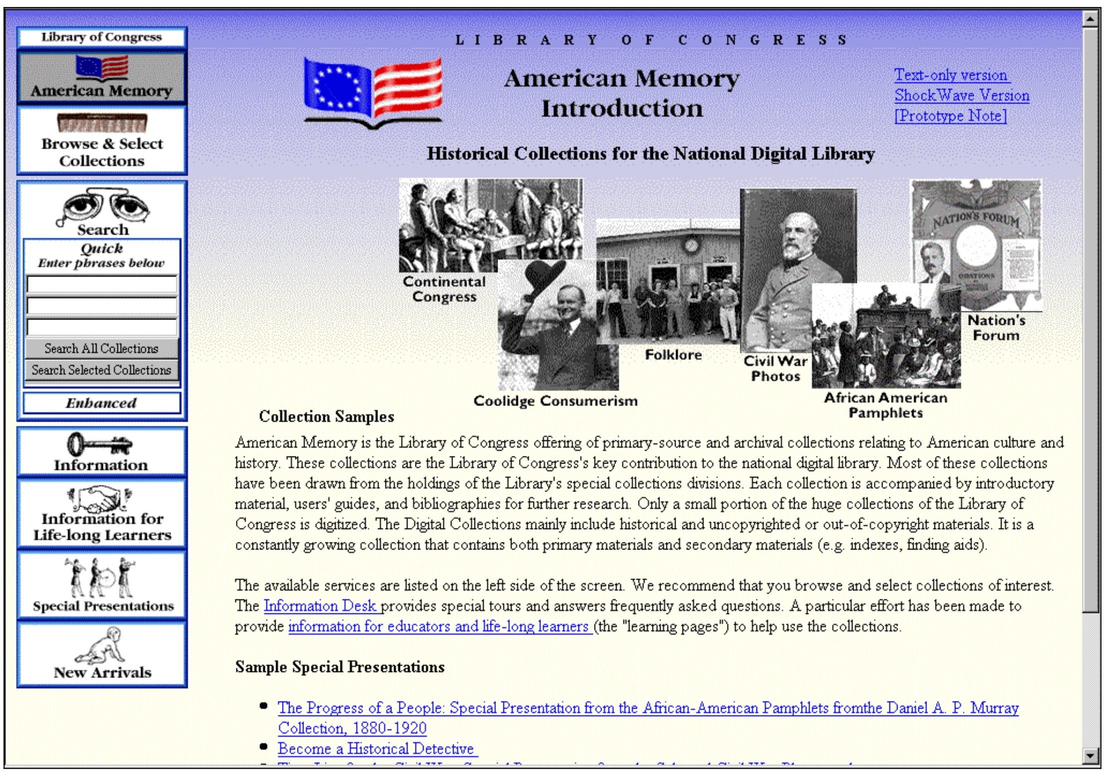

Figure 1a: Opening screen design for the American Memory Collections, showing animated samples of the collections and a quick search available from the home page.

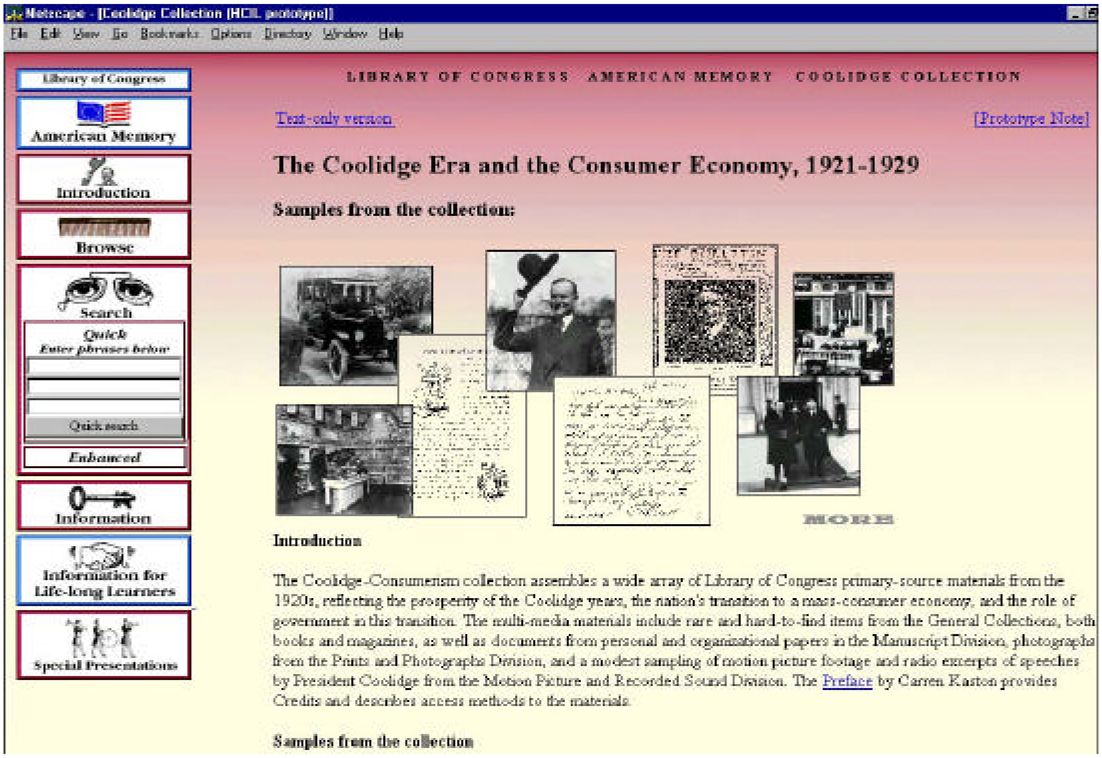

Figure 1b: A consistent layout for the Coolidge Collection

Introduction and Samples. On the introduction page users see animated samples from the corresponding collection(s), a short textual introduction, and links to sample special presentations written by experts. The clickable samples were meant to provide immediate access to some of the treasures of the collections instead of forcing users to navigate through deep menu hierarchies or execute a query. Users often have a vague idea of what they hope to find. Presenting many samples immediately at the home page provides an enjoyable experience to first time users and "web surfers", and may help users formulate their search by progressively gaining a subjective appreciation for the type of materials available. Expert users can of course ignore these samples and move immediately to the objects of interest or pose queries using the search feature.

Several different ways to present samples were explored. Sample presentation can be entirely under user control (as in the AM interface) or automatically animated (as shown in the CC). The selection of samples from the collection could be automated, but the collection curator or other LC staff member will of course do a much better job. An ideal sample selection will be both attractive and representative. It may be feasible and popular to periodically solicit sample sets from groups through educational contests or to develop thematic samples based on current events or historical celebrations.

Browse. For both interfaces (AM and CC) we developed novel, Java-based browsing interfaces. For the American Memory interface we developed a Collections Browser applying dynamic query principles to the exploration, filtering and selection of collections. For the Coolidge Collection we developed a generic collection browser (called WebTOC), especially designed to be usable for unprocessed collections as well as processed collections.

These browsers provide overviews to large sets of collections or sets of items and give users a sense of the amount of material and the attributes of these materials. Overviews help in selecting collections and locating sets of items in a collection with similar attributes. Such tools also help users learn about the collections in general, what topics are covered, and what kinds of materials can be found, thus helping them to form expectations and formulate search queries.

<u>Collection browser</u>. The Collection Browser uses dynamic queries [Ahlberg, et al., 1992]. Dynamic queries extend the "direct manipulation" style of interaction to search

tasks by closely coupling search specifications with the display of results. This interface (see Figure 2) features a visual overview of all the collections (a zoomable timeline with a bar for each collection showing its time coverage), a visual representation of a query using a collection of widgets (three menus for location, topic and format) and tight coupling between all components of the display (timeline, filters, list of collections and short descriptions--all updated concurrently). Users browse the database by interacting with the interface widgets (filters, zooming of timeline or simple cursor selection). Each change is in effect a new query, the results of which are immediately and continuously shown in the display. These dynamic queries support progressive refinement of the search, continuous reformulation of goals, and visual scanning to identify results.

The collection browser requires consistent metadata for all the collections. The first iterations of this design uncovered problems with the existing metadata. Often metadata was not available for the materials on the item level or was inconsistent across collections, making it impossible to develop such a browsing interface at the item level across all collections. The LC/HCIL team decided to define a set of consistent attributes at the collection level to support across-collection browsing. The LC team members defined the metadata format, and metadata were generated as necessary for the collections. The timeline, collection list and menus are created automatically from the metadata provided.

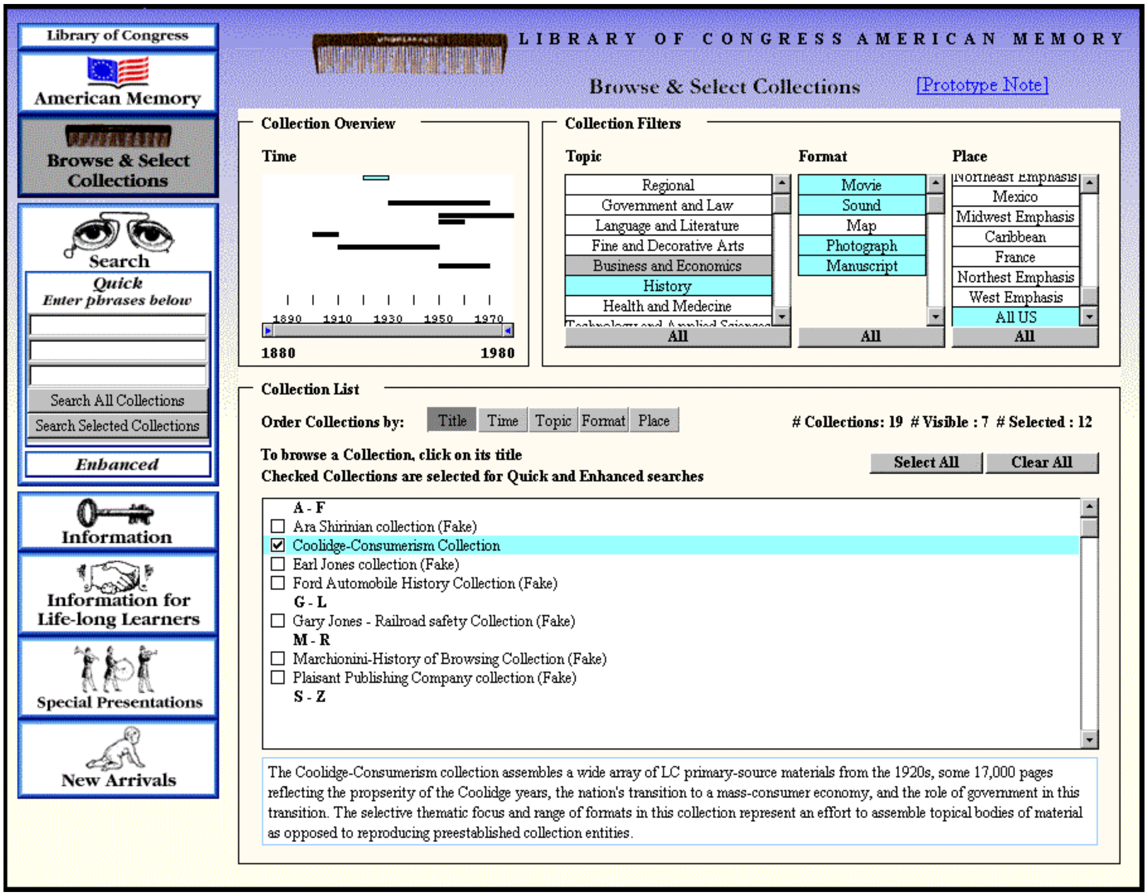

Figure 2. Overview and browsing of collections. At first all collections are shown on the timeline and in the list. Here the user filtered the collections by selecting the topic “Business and economics”. The Coolidge Collection is highlighted revealing its attributes and short description at the bottom of the screen

The principles of the collection browser could also be used to browse items of collections if consistent metadata is available for each item of the collection. In specific collections with homogeneous data types, this may be possible, but this is not the case for most large, diverse DL collections such as American Memory. In fact, one of the reasons the

Coolidge Collection was selected for study was for its multimedia content. Therefore we had to develop a more generic technique which could be apply to any set of digitized items.

<u>WebTOC for browsing an individual collection.</u> The individual collection browser tool is called WebTOC [Nation et al, 1997]. One difficulty webmasters and users alike have is representing the quantity of information and its distribution within a set of linked documents. Digital librarians must manage huge directories and need tools to visualize directory structures and data types. Knowledge of the quantity and types of documents available can be helpful to users as they decide whether a web site may be interesting or useful. To browse large numbers of documents in the growing collections is not an easy task since the size and extent of such large collections makes them hard to understand. WebTOC is a tool that can help both end users browse and library staff develop and organize the collections.

WebTOC automatically generates a hierarchical table of contents of a site (or, in the case of the NDL, a collection) using two different strategies: following existing links or using the underlying directory and file structure. Following links is appropriate for existing web sites, while using the directory structure is appropriate for newly digitized "unprocessed" collections, which have not yet been linked, indexed or annotated. WebTOC consists of two parts: a parser program to generate a table of contents (TOC) file representing the site, and a user interface (Java applet) to display the TOC and allow interaction within a conventional web browser (Figure 3).

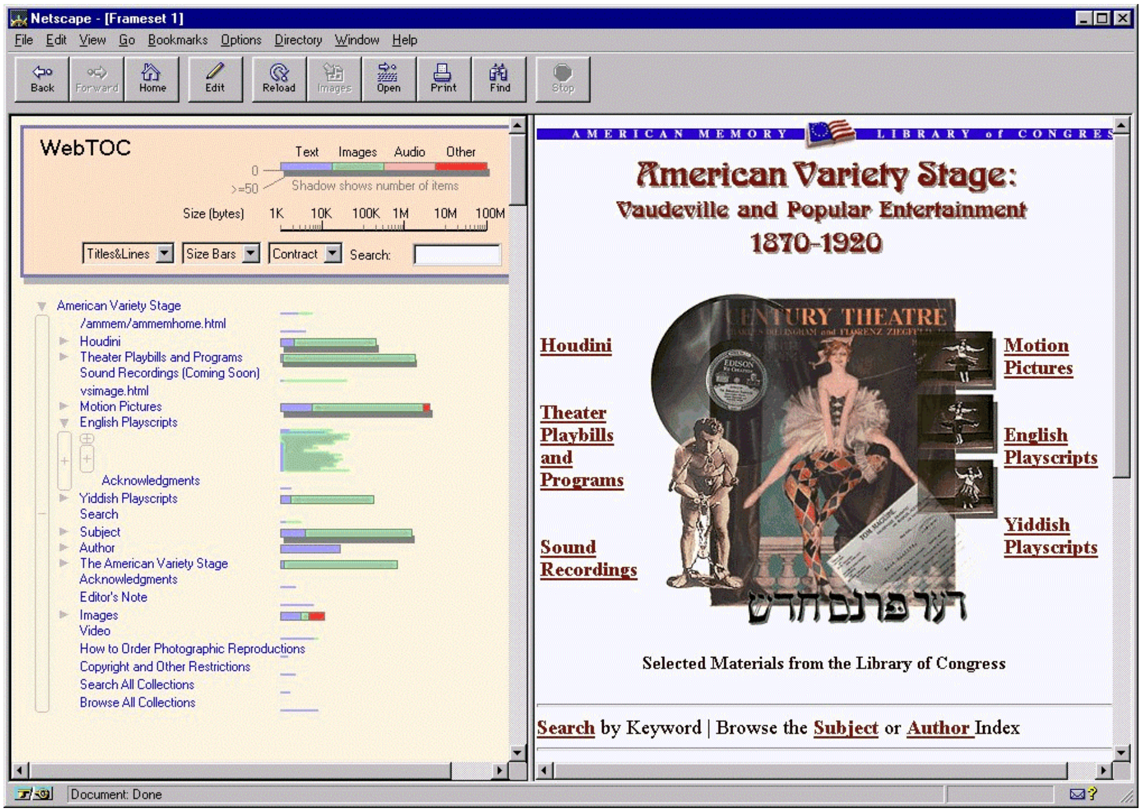

Figure 3: WebTOC shows on the left the table of contents of the collection/site "American Variety Stage”". The top portion of the WebTOC frame is a legend and control panel for WebTOC. Links are listed with a bar that represent the volume of information available when following that link. Users can then expand the hierarchy (e.g. here "English Playscripts" has been expanded and the labels removed to compare file sizes). Clicking on any of the labels or individual lines causes that file to be displayed in the frame on the right.

A two-frame layout allows users to expand or contract the table of contents presented in an outline format on the left frame, while inspecting any selected item in the right frame. As summarized in the WebTOC legend (top left of Figure 3) individual pages are shown with individual lines while bars aggregate groups of files behind a link or directory to represent the total size of the included documents. Color is used to represent file type, length the overall size and the shadow below the bar is proportional to the number of documents included. The number of documents and size represented are displayed in the browser status area when the user's cursor is over the bar. Figure 3's TOC shows that "Sound recording" is very small (close to empty) in comparison to the other categories like "motion pictures", an observation that would have been impossible to guess from the home page on the right. "Yiddish playscripts" is also smaller than "Motion pictures" but the shadow is also very small suggesting that there are few pages of playscripts. "English playscripts" has been opened using the arrow button on the left of the name and the individual lines show individual pages. WebTOC also provides a text search capability for the table of contents and displays results in the context of the table of contexts hierarchy.

WebTOC was useful to us to as we worked with the unprocessed Coolidge collection and should be a useful tool for library staff preparing or reviewing new collections. It is also likely to be useful for the general public as it provides an alternative access to the materials. It is particularly useful to explore a collection when searching is inadequate because of the lack of metadata, finding aids, or lack of any searchable text; or because users are not able to specify their search appropriately. Additionally, it provides a way for expert users to quickly jump to objects deep in a hierarchy without necessarily following a series of explicit links on pages that each take time to load.

A pilot study was conducted to evaluate the usefulness of WebTOC versus a plain browser for Web site navigation tasks. Results suggest that the complexity WebTOC adds to the display makes simpler tasks more difficult but in the case of more complex tasks it adds enough value to outweigh display complexity. Users found WebTOC easy to learn and its presence also increased their "organizational feel" for the site. The full paper with detailed results is available at http://otal.umd.edu/SHORE.

Search. Because the NDL will serve a broad range of users from web novices to professional search intermediaries, several search widgets were included. Quick Search is available as part of the anchoring toolbar on the left (Figure 1). Because search is so widely available in web interfaces, users will expect to search the entire American Memory collection by typing a few words or phrases. In the AM interface two quick search buttons are provided: one to search all collections, the other to search only the collections selected in the "Browse and select" page (Figure 2). In the Coolidge collection, the quick search (for the purposes of the prototype, a simple string search was applied, LC used the Inquery engine for ranked retrieval during the period this work took place) applies to the Coolidge collection only.

Enhanced Search supports Boolean queries and the restriction of the scope of the query to collection objects of specified formats (e.g. photograph collections). It also gives access to all the attributes common to the collections searched. This is particularly useful in specific collections where metadata is available but not applicable for restricting the query across all collections since common data (e.g., year ranges or geographic locations) does not exist for all items. If, for example, users have restricted the scope of the query to three photograph collections which all have consistent data about the type of photograph and whether the item is in color or black/white, it is feasible to automatically detect those common attributes and dynamically include this as a filter on the enhanced search screen. The Enhanced Search screen is organized into four blocks: scope, attributes, result, and variants. Scope allows users to define where to search, attributes let users enter text (what to search on) and define fields, results allows users to set the number of hits and set sort settings, and variants give help to users on formulating the query (Shneiderman, Byrd, & Croft 1997). In the toolbar of the left side of the screen, the search tab is a rather complex box with entry fields and multiple buttons. A click on the background of the tab, or on the word "search" itself brings a help page for the search functions.

Display of results. Displaying results is one of the most underdeveloped aspects of web-based searching. Professional intermediaries searching commercial systems such as Dialog or Lexis are provided with a variety of print and display options but web users must typically deal exclusively with long lists of ranked title lists. In multimedia digital libraries this approach to results display is particularly problematic as objects of different data types all look alike. In the NDL, result lists are composed of heterogeneous materials in both format (photographs, manuscripts, motion pictures, letters from a person, one issue of a periodical, etc) and by the granularity of items (collections, items, finding aids, etc.). To address this problem, different results display options were created to distinguish different levels of representation and provide some sense of the volume of material returned. This was done in some rather rudimentary way in the AM interface and in more detail in the CC interface.

In both cases the screen has three tables: a summary of the query, a summary of the results with links to format-specific preview tools, and a detailed results list. Showing the user the query is another example of conceptually anchoring them in context. Links to the specific preview tools is an added value service that operates more smoothly than the few tools on the web that typically launch special viewers or plug-ins.

In the CC the results screen illustrates that search results should indicate what types of object were returned by the search: a whole collection, a simple item, a finding aid leading to many items, or a finding aid for a non-digitized collection (Figure 4). For this prototype, an icon was adopted for each data type. The most significant aspect of this design is that the result list not only shows the objects returned by the search (e.g. a transcript or a bibliographic record) but also gives links to all the other representations of that object. This added value service has great potential for facilitating browsing. For example, if the search word was found in the table of contents of a manuscript, the result list provides a link to the table of content but also to the bibliographic record of the manuscript, its transcript or related expert note, and any other representations of the manuscript. This provides a much richer results set, allows supplementary filtering or ordering of the result lists, and possibly reduces the navigational effort because users could decide to jump directly to the representation of greatest interest. Of course this interface relies on the existence of underlying links through the different representations of given item. Our study of the Coolidge collection seems to indicate that this should be possible to automate as the naming scheme of the files was enough for us to guess the type of representation to be found in a given file and to link to other related files.

One of the aspects not addressed is making visible an estimation of the number of items NOT covered by the search engine (because of lack of metadata or any textual information). In the case of the NDL this could be a large proportion. To alert users to this fact, a warning message could be displayed while the search is performed, and recommending the use of browsing techniques.

Finally we illustrate in WebTOC that results could also be displayed in the context of the collection table of contents. When a search is done is WebTOC the results are shown by opening all (and only) the tree branches leading to result hits - therefore showing results in context. Figure 5 illustrates this feature.

These techniques for handling results are interesting beginnings for what we believe is an important but typically ignored component of digital library services. Significantly more research and development work is needed to find ways to browse and use search results.

Format-Specific Preview Tools

Both search and WebTOC may yield long lists of items, which can be tedious to peruse. Browsing through a collection of 500 photographs using WebTOC or a search result list will take at least 500 operations (i.e. click-and-wait); worse yet, browsing through a collection of 30 videos can takes hours of download time and megabytes of disk space.

To illustrate that it is possible to speed the browsing of large numbers of items with specialized tools, an image browser, a video browser, and a page browser were developed to allow users to peruse previews of the items before requesting details. These tools have been developed in the Coolidge Collection prototype but are not specific to this collection. They can be accessed from the result summary of the result screen (e.g. the summary indicates the total number of photos and provides a link to the image browser loaded with those images), or from any collection overview (WebTOC, topic index, table of contents, etc.).

Image browser. The ability to search or query a database of images is a necessary function for digital libraries containing images. The query feature is needed in order to accommodate users who wish to find a particular image from among the many images present in the collection. Upon retrieval additional functionality is required so that users can quickly and effectively browse the retrieved set of images.

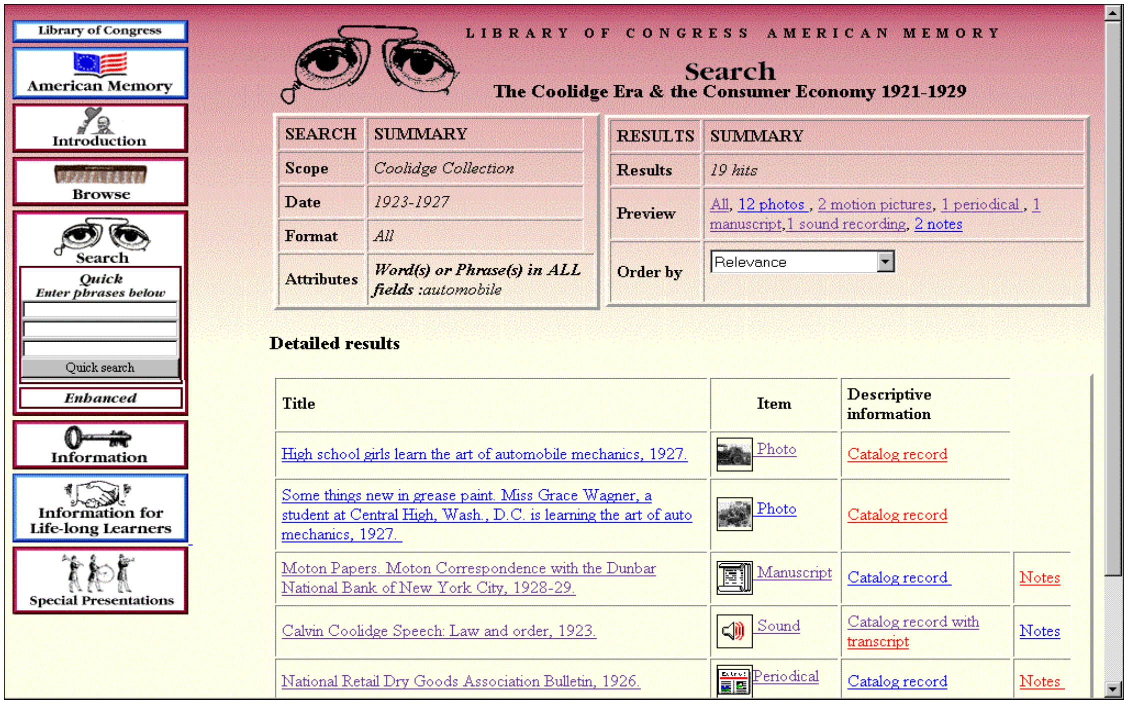

Figure 4: Search results display screen showing the search summary, result summary and detailed list. Each item "hit" by the search (e.g. a sound transcript) is displayed along with the item itself and other surrogates for the item (catalog record or notes).

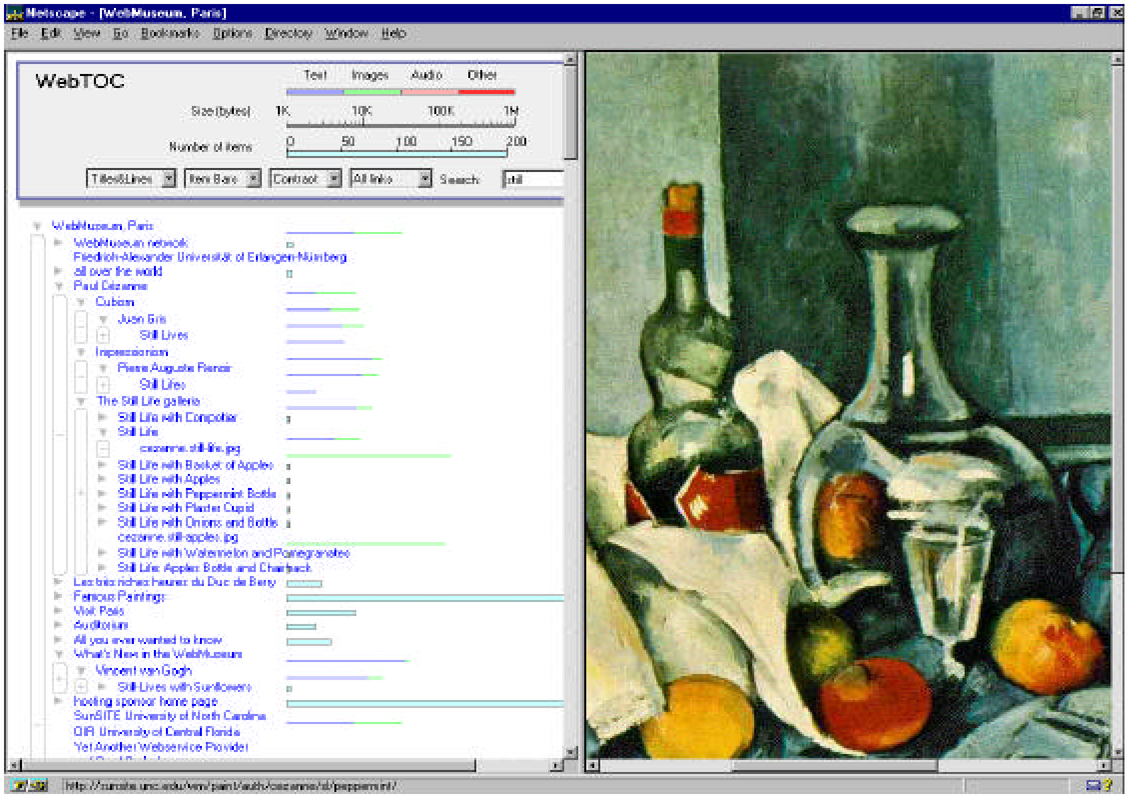

Figure 5. WebTOC showing in the table of contents the results of a search on: "still" in the context of the hierarchy of the "WebMuseum Paris" site.

One common method used for querying image databases is to allow the user to enter a text description of a target image. The system uses this text description to match the query with text descriptions of images in the database and outputs the possible matches. This method requires linguistic indexing for visual content and even with the best indexing is subject to all the limitations of Boolean or free text searching. Alternative methods of searching an image database have arisen in order to augment or replace text descriptions of images. One method is Spatial-Query-By-Sketch (Egenhofer, 1996). Spatial-Query-By-Sketch requires the user to draw a simple sketch of the image sought. It then executes an algorithm to generate matches for the query based on the input of the user sketch. While this method does provide the user with a more direct model for searching, the user will still eventually be left with a series of possible matches, some of which are incorrect translations of his/her sketch. Other techniques use color matching(e.g. IBM Query-by-Image-Content [Flickner et al., 1995], UC Berkeley DL Project {Forsyth, Malik, & Wilensky, 1997]) or icon-selection languages (e.g. Nishiyama, Kim, Yokoyama, and Matsushita, 1994) or combinations (e.g. Chuah, Roth, Kolojejchick, Mattis, and Juarez, 1995). In all cases the object sets retrieved must be previewed and browsed.

The need for a well-designed and efficient image browser is apparent. If a query is not performed, the image browser is the only method of searching a given database. If a user does perform a query, the user must then browse through the possible matches returned by the query. Also, query algorithms can sometimes eliminate (and not display) images for which the user may be searching, a browser is needed in this case in order to find such images.

The image browser created in this project can display 10, 20 or 50 thumbnail previews of images at a time (Figure 6 a, b and c). The top frame shows the thumbnails of images, summary information of the search results and user controls. The center frame can be used to view a slightly larger version of the image and the bibliographic information (Figure 6b) or the full resolution image (Figure 6a). Selected images shown with their flag up are collected in the bottom frame as thumbnails and can be saved on-the-fly in a personal web page.

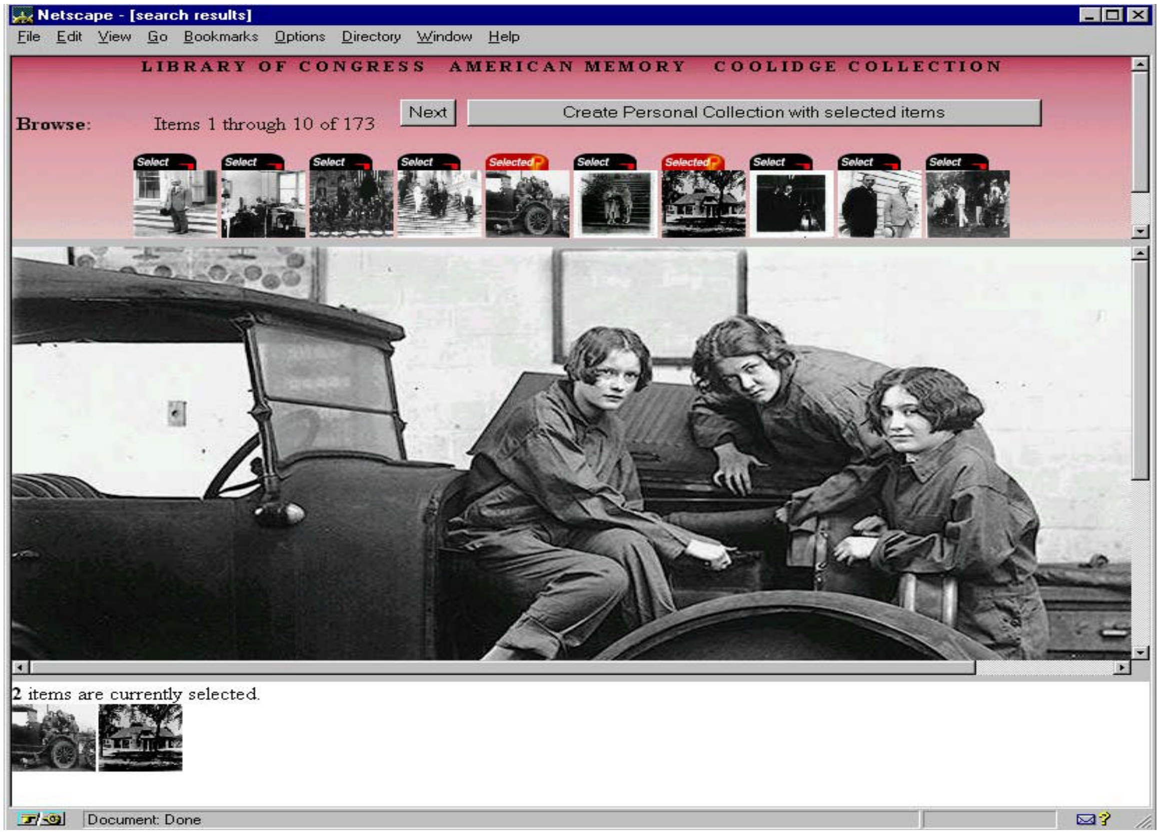

Figure 6a: The image browser can be used to view results or possibly entire medium size collections. Here 10 thumbnails are shown at once. The center frame is used to view the full resolution image. Selected (flagged) images are collected in the bottom frame as thumbnails and can be saved in a personal web page.

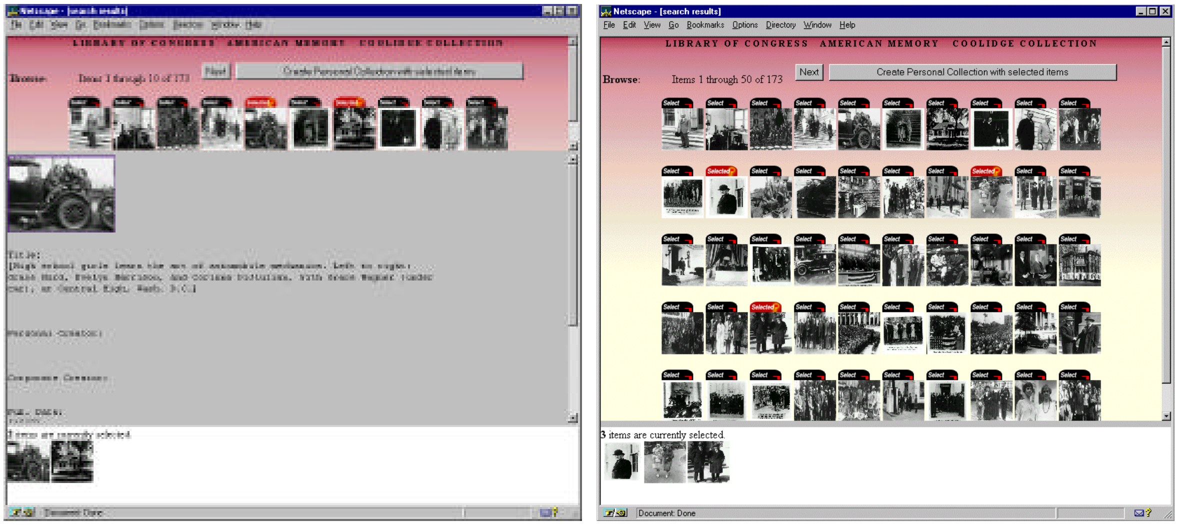

Figure 6b and 6c: In 6b the user chose to view a larger thumbnail version of the image and the descriptive textual information . In 6c, 50 images are shown at once.

The number of image per screen might be a function of the user's preference or screen resolution. With 50 per page it is possible to peruse 500 images in only 10 operations - and 10 screens (Figure 4c). Of course users will still request details about some of the items (bibliographic record or full image) but the browser provides a good way to evaluate the set of result items or even collections as a whole. It complements nicely the textual lists often available. Note that the image browser does not require that any descriptive material be available about the images, and in some ways is an electronic equivalent of the flipping through boxes of undescribed materials at the library (e.g. in the Prints and Photograph Reading Room.)

Visually browsing large sets of thumbnails of images supports users better than browsing a textual list of titles. Evaluating visual information with only textual descriptions makes selection harder for users, especially since image attributes like colors, light, style, are important for discrimination and can often only be judged by actually looking at the picture. Thumbnails make efficient use of the screen real estate and provide enough detail for the users to make selections. By compressing the visual preview and detailed view on one screen this browser reduces navigation for users and saves time compared to the present interface where users have to go back and forth between the list of titles and the individual images.

Gathering Materials. After finding materials of interest users are likely to want to gather the materials in their own pages. Students will use pictures or manuscripts in their reports or point to expert notes. Teachers will want to create pages listing materials of interest for their classes. Bookmarks can be used but large numbers of bookmarks are hard to manage and cannot be shared easily. HTML editors can be used but require training to be used. A gathering tool was created using JavaScript to assist users in taking full advantage of the invaluable materials by saving them and using them later for their own purposes (see Figure 7). LC staff to quickly generate aggregations of items for different purposes, for example, to create a Hot Topics page can also use this tool. The prototype only includes the gathering of photos but should be extended to other types of materials. When users click on the "Create personal collection" button in the image browser, a new HTML page is created on the fly. It is a generic page which mainly consist of a table with an image on the left and a generic text on the right. If the user has an HTML editor, the generic page can be saved on users' local disk space and edited rapidly. Users can add text or annotate the images and use them.

Figure 7: Gathering tool to create personal collections. A template is generated automatically, which can be saved locally for editing, or used as-is immediately in a class or work presentation**.**

Video browser. Digitized motion picture documents are important components of some digital library collections. Users can browse these collections or search on the textual descriptions and in some cases search visual attributes such as color, optical flow, luminosity, and shape. When browsing a collection or a result set from a search it is very important to provide suitable surrogates for the videos, since the files of these documents are usually very large, even if the clips are short.

The video browser is designed to review sets of video rapidly (Figure 8). A list of videos is shown with the first or most representative frame of each video. When a video is selected users can see a preview made of selected frames of the video. Several research groups are working on techniques to extract representative or key frames from videos (e.g. Wactlar, 1996). For this prototype, an algorithm developed at the University of Maryland Center for Automation Research (Kobla et al. 1996) was used to parse MPEG-compressed video sequences into shots based on their physical structure and further into scenes by identifying changes in content and camera motion. Alternative mechanisms to present those key frames including animation, image array, and multiple videos were incorporated.

Figure 8. Video browser with animated display of key frames. In the left frame, one key frame from each motion picture represents the respective films in the set (a film collection or here a search result with 2 films). The user can click on these

to select the film to preview. On the right appear the bibliographic information and an animated preview of the film.

A number of user studies were conducted to explore video key frame control mechanisms. One study examined different key frame slide show display rates and found that key frame rates between 8 and 12 fps potentially define a functional limit in object identification accuracy [Ding and Marchionini, 1997]. This effect was found to be less pronounced for gist extraction purposes. Another study [Slaughter et. al., 1997] examined the efficacy of multiple key frame slide shows running concurrently. Results suggest potential for two or three concurrent displays but very poor performance for four displays. Other studies are currently underway.

Page Browser. Another type of browser that can have an important impact on the usability of collections such as American Memory is a page-oriented browser. Typical interfaces provide a table of contents with links to the individual pages, requiring the user to go back and forth between the index page and the page images to go from one image to the next.

In some ways, digitized pages are merely images and the image browser could be used for manuscripts or journals. But we found that a better browser could be designed for this important category of materials, which have some kind of natural page ordering, and an often-consistent page size. Our study of the Coolidge collection suggested that this is the case for journals, book chapters, catalogs, and often for manuscripts. First a browser was manually prepared for a journal, a manuscript and a book chapter, then a tool that partially automates the production of such browsers was created.

The basic design consists of three frames (Figure 9). On the top frame the title of the collection and journal or manuscript name appears, with links below it to any existing complementary materials (bibliographic record, notes, etc.) and to a screen of thumbnails of the pages. On the left frame an overview of the journal or manuscript is shown. It can consist merely of a list of page numbers, a table of contents or a richer finding aid. On the right frame the actual pages are shown when selected in the overview. The table of contents is tightly coupled with the page image viewer frame, and updates as the user changes pages to continuously show the location of the page the user is looking at in the document (with an arrow).

These browsers can be accessed from the result list when the item is actually a collection of pages, or from any collection overview which has clusters of pages (WebTOC, indexes etc.). This browser uses mainly HTML frames (and CGI scripts for the arrow).

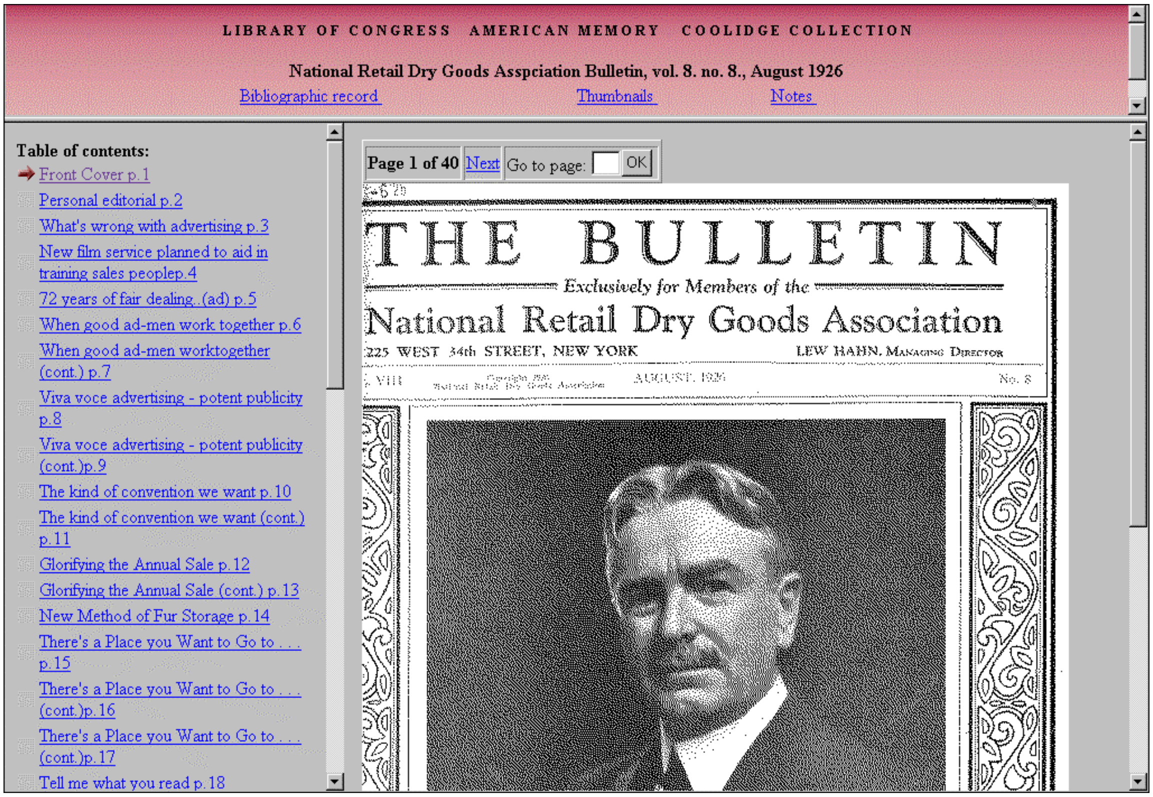

Figure 9. Page image browser for periodicals. The top of the screen shows the title and number of the serial, the collection it belongs to and other information pertinent to the document as a whole. The user can select whether to look at the pages in the thumbnail view, full-page images, or full bibliographic records. On the left a list of the pages in the document appears. An arrow indicates the page being shown.

One of the main goals was to provide overviews of the page collections, to give a feeling of the type (content previews; text or images; typed or handwritten, text + handwritten notes, etc.) and length (no. of pages, size of pages, etc.) of the document. The table of contents provides textual overviews, while visual overviews are provided by the thumbnail view. To reduce feeling lost and to give better orientation to users in the document the tightly coupled table of contents and page image displays were implemented, always giving the user a feeling of where in the document the page is located in relation to other pages. Contextualization and orientation may be more important in the case of books and serials where the structure of the document itself adds to understanding the content.

Our study of the Coolidge collection and the work of others (e.g. [Adar et al 1995, Noreault and Crook 1995] suggests that it is possible to automate the production of such browsers. Once pages are digitized, aside from writing notes and documentation (which must be done by humans anyway), creating the online HTML documents for presentation of the material is a highly repetitive and tedious process. A tool that automates the process such that minimal human intervention is needed was developed. It may be used for manuscripts, journals, book chapters or catalogs (Figure 10a and b)

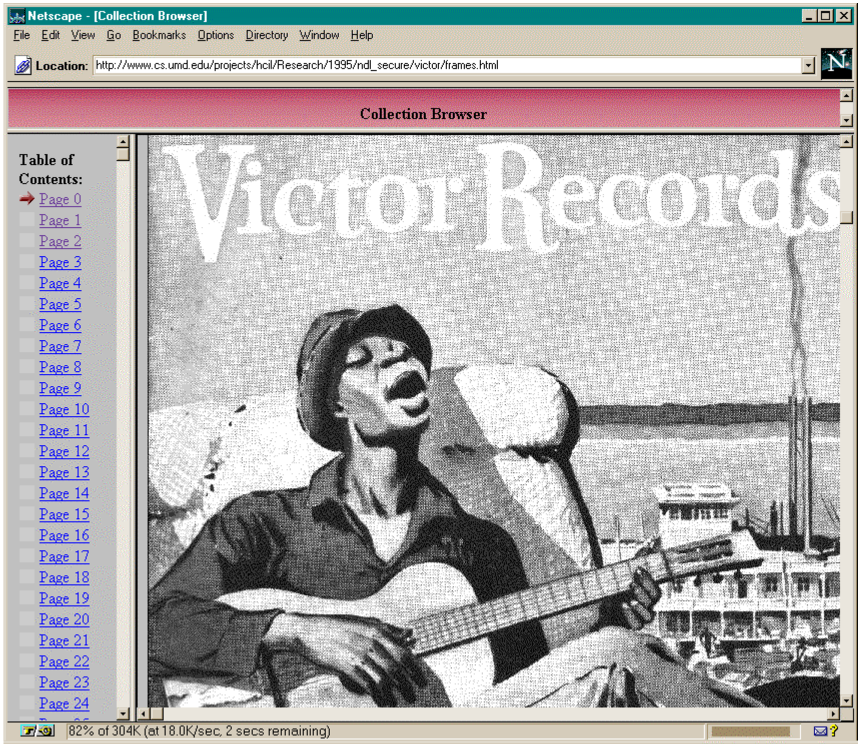

Figure 10a. Example of rudimentary page image browser created entirely automatically. This kind of limited browser can be accessed from query results or WebTOC to browse unprocessed collections. Simple manual edits can lead to improved versions such the one shown in Figure 10b

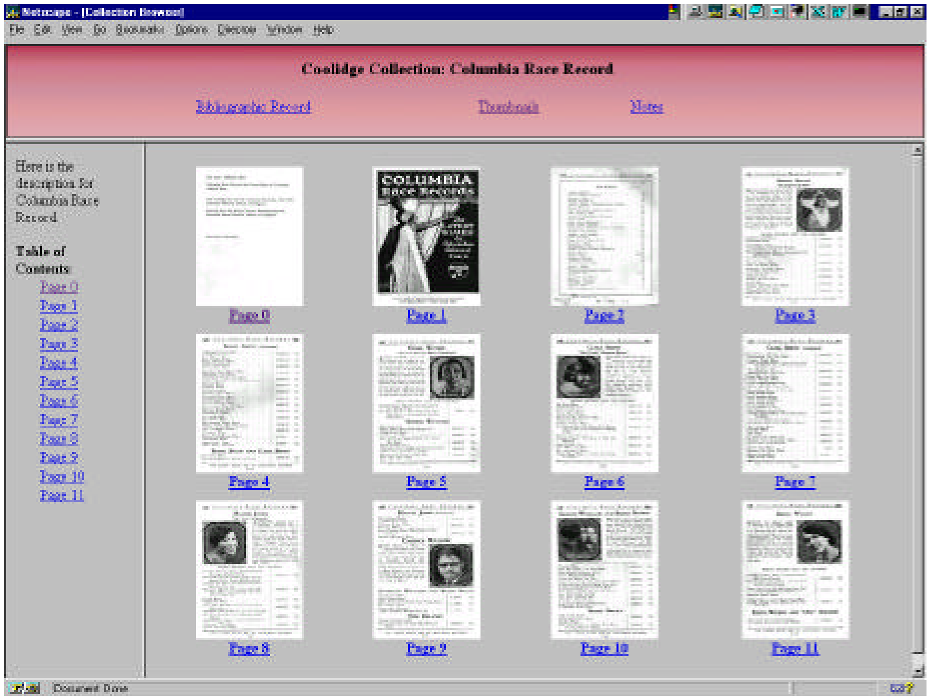

Figure 10b: Another example of automatically built browser. The table of content and thumbnails were generated automatically. The bibliographical record would be created or located manually, providing the title and short description to be extracted automatically from the record. Notes will always be manually generated but can be linked automatically when available.

Discussion and Future Research

This project has demonstrated interface techniques that bring primary information resources to users quickly and easily, anchor users in a consistent context to minimize disorientation, and support rapid relevance decisions through overviews and previews. These interfaces and tools may be used with digital collections that have consistent metadata, as well as techniques that provide users with good access even though metadata is absent or inconsistent. These distinct designs illustrate techniques that may be applied to highly homogeneous collections or across the more typical uncataloged or nonstandardized collections found in digital libraries. The designs illustrate solutions to the challenges inherent in DLs: serving wide ranges of users and tasks with consistent interfaces that minimize disorientation and provide good overviews and preview facilities. Aspects of this prototype (e.g., vertical table of contents, sample items on the first page) have already been incorporated into the most recent version of the LC NDL and American Memory site but user testing is needed for the more complex interfaces.

Digital library technology will develop rapidly in the coming years and the Library of Congress continues to lead the way with innovative and effective interfaces for its NDL Program. This project aimed to define a trajectory of interface development that not only facilitates research-based designs for today's technology but also will assist the Library in shaping future iterations that serve broad communities of users. A number of research and development issues related to DLs in general emerged as a result of this work.

First, principles and guidelines for user-centered, iterative design of DLs must be developed and tested. The first step in this direction is for DL staff to adopt user-centered design perspectives that include regular user needs assessment, ongoing usability testing, and iterative design procedures. DLs are not static entities. Just as physical libraries are constantly evolving, DLs will grow and change with users and technology.

Second, DLs must serve diverse user communities and will need to develop appropriate interfaces for varied users and needs. Rather than depending on users or their agents to deal with a fixed interface, we encourage alternative interface suites that users may select according to their specific experiences and immediate needs. Although we sketched a virtual information desk component in our phase one work, we focused on building specific browsing tools rather than prototyping the various interface alternatives. This is important follow up work.

Third, the development of the various tools and our own struggles with managing megabytes of data highlighted the need for digital librarian toolkits. Integrated tools that digital librarians use to manage materials and web sites and to communicate with users are sorely needed. The content management tools will help librarians in acquiring materials for the digital library and presenting them to the users, that is, to collect, store, index, preserve, display, create, and secure content. Tools to automatically test codes for different platforms; find inconsistencies in words, icons, and widget suites; and manage different versions must be invented and added to these toolkits to support creation of the interfaces. Maintenance tools such as transaction log analyzers, site map generators, and automatic FAQ updaters will also be useful to digital librarians who create, move, update, and weed digital collections and serve diverse user populations. Many tools in the content management category exist as standalone entities for building and maintaining web sites (e.g., file format converters (e.g., JPEG to GIF), link checkers (e.g., robots to find broken links), automatic surrogate extractors (key frame extractors, text summarizers), parsers (e.g., find SGML or other markup objects), and encryption routines (e.g., watermarks, cryptolopes). Tools needed in the communication category include tools to support user services such as reference, to create and maintain user profiles, to collect and analyze transaction logs, and to build and maintain community participation. Digital librarians and web masters already acquire and develop many of these capabilities as standalone applications or workarounds. What is needed are unified toolkits that provide such tools covering all areas of digital library work via a common interface.

Finally, DLs are rooted in organizations; organizational culture will influence the development and maintenance of the DL. It takes considerable time and effort to understand the various needs of different stakeholders in complex organizations. For example, at LC, although representatives from various reading rooms and divisions were world authorities in their specific collections and work areas, no one person fully appreciated the work practices, user needs, metadata requirements, and existing website investments of all LC units. The many meetings that took place during the first phase of the project served to bridge these gaps and develop a more global set of problem specifications from which design could flow. All DLs will face similar challenges and participants should be prepared for the time investments related to the social and political dimensions of DL development. Additionally, we should prepare early on to assess the inevitable impact the DL will have on the hosting organization. At LC, for example, the attention given to K-12 users of the NDL may eventually influence the overall mission of the Library and the work climate of the physical library itself.

Acknowledgements

The Library of Congress provided partial support for this project. We want to thank David Nation who developed WebTOC while visiting HCIL from the Department of Defense, and Tom Bruns (Collection Browser), Ara Shirinian (image browser and video browser), Steve Karassik (result display), Teresa Cronnell (graphic layout), and Ben Shneiderman (feedback over course of project). We also want to acknowledge the cooperation of the team at the Library of Congress headed by Laura Campbell (Martha Anderson, Maryle Ashley, Ardie Bausenbach, Mike Black, Jeff Bridgers, Beth Davis-Brown, Ernie Emrich, Floris Flam, Georgia Higley, Lois Korzendorfer, Sandy Lawson, Betsy Mangan, Gary Fitzpatrick, Barbara Natanson, Virginia Sorkin, Susan Veccia, PeterYeager and Walter Zvonchenko). The final prototype is available on the project web site http://www.cs.umd.edu/projects/hcil/Research/1995/ndl.html

References

Ahlberg, C., Shneiderman, B., Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays, *ACM CHI '94 Conference Proc*. (Boston, MA, April 24-28, 1992) 313-317.

Adar, E., Hylton, J. (1995). On-the-fly Hyperlink Creation for Page Images. *Proceedings of the Second Annual Conference on the Theory and Practice of Digital Libraries (Digital Libraries '95)* (pp.173-176).

Bederson, B. & Hollan, J. (1994). Pad++: A zooming graphical interface for exploring alternative interface physics. *Proceedings of Seventh Annual Symposium on User Interface Software and Technology* (Marina del Rey, CA, Nov. 2-4, 1994). NY: ACM Press.

Brown, M, and Shillner, R., The deckscape Web browser, *ACM CHI '96 video*, Summary in *CHI' 96 Companion*, 418-419, ACM New York.

CACM 95, Special issue on Digital Libraries, *Communications of the ACM*, April 1995, **4**, 29-39, ACM New York.

Card, S., Robertson, G., York, W., The WebBook and the Web Forager: An Information Workspace for the World-Wide Web, in *CHI'96 Proceedings*, ACM, New York, 1996.

Chuah, M.C., Roth, S.F., Kolojejchick, J., Mattis, J. & Juarez, O. (1995). SageBook: Searching Data-Graphics by Content. *Proceedings of the Conference on Human Factors in Computing Systems(CHI '95)* (pp. 338-345). New York: ACM.

Conklin, J. (1987). Hypertext: an Introduction and Survey, *IEEE Computer*, 20(9), 17-41.

Doan, K., Plaisant, C., Shneiderman, B., Query previews in networked information systems, *Proc. of the Third Forum on Research and Technology Advances in Digital Libraries, ADL '96* (Washington, DC, May 13-15, 1996) IEEE CS Press.

Ding, W. and Marchionini, G. (1997). A Study on Video Browsing Strategies. *Technical Report* CS-TR-3790, UMIACS-TR-97-40,CLIS-TR-97-06, University of Maryland. http://www.learn.umd.edu/wp/speedexp.html

Egenhofer, M., (1996). Spatial-Query-By-Sketch. Proceedings*: IEEE Symposium on Visual Languages*, 60-67.

Flickner, M., Sawhney, H., Niblack, W., Ashley, J. Huang, Q., Dom, B, Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D., & Yaker, P. (1995). Query by image and video content: The QBIC system. Computer, 28(9), 23-31.

Forsyth, D., Malik, J., & Wilensky, R. (1997). Searching for digital pictures. Scientific American, 276(6), 88-94.

Greene, S., Marchionini, G., Plaisant, C., & Shneiderman, B. (1997). Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *Technical Report* CS-TR-3838, UMIACS-TR-97-73, University of Maryland.

Hearst, M., Tilebars: Visualization of term distribution information in full text information access. In *Proceedings of CHI'95,* ACM New York, 1995. 59-66.

Heath, L., Hix, D., Nowell, L., Wake, W., Averboch, G., Labow, E., Guyer, S. Brueni, D., France, R., Dalal, K., & Fox, E. (1995). Envision: a user centered database of the computer science literature. *Communications of the ACM*, 38(4), pp. 52-53

Kandogan, E., and Shneiderman, B., Elastic Windows: A Hierarchical Multi-Window World-Wide Web Browser, *University of Maryland Technical Report* CS-TR-3789, ISR-TR-97-56 (May 1996)

Kobla, V., Doermann, D., & Rosenfeld, A. (1996). Compressed domain video segmentation. Technical Report CAR-TR-839 CS-TR-3688, University of Maryland.

Landauer, T. & Nachbar, D. (1985). Selection for alphabetic and numeric menu trees using a touch screen: Breadth, depth, and width. Proceedings of CHI '95, Human Facotrs in Computing Systems, NY: ACM Press. 73-78.

Marchionini, G. (in press). Digital Library Research and Development. *Encyclopedia of Library and Information Science* (A. Kent, Ed.). http://www.glue.umd.edu/~march/digital_library_R_and_D.html

Marchionini, G., Plaisant, C., Komlodi, A. (1996). User needs assessment for the Library of Congress National Digital Library. *University of Maryland Technical* Report CS-TR-3640, CAR-TR-829, CLIS-96-01.

Nation, D.A., Plaisant, C., Marchionini, G., and Komlodi, A. (1997) Visualizing web sites using a hierarchical table of contents browser: WebTOC. *Proceedings of the 3rd Conference on Human Factors and the Web*, 1997.

Nielsen, J. (1990). *Hypertext and Hypermedia*. Boston: Academic Press.

Nishiyama, H., Kin, S., Yokoyama, T., & Matsushita, Y. (1994). An Image Retrieval System Considering Subjective Perception. *Proceedings of the Conference on Human Factors in Computing Systems(CHI '94)* (pp. 338-345). New York: ACM.

Noreault, T.R., Crook, M.A, Page Image and SGML: Alternatives for the Digital Library. *Proceedings of the International Symposium on Digital Libraries 1995*.145-150.

Plaisant, C., Marchionini, G., Bruns, T., Komlodi, A., Campbell, L. (1997) Bringing Treasures to the Surface: Iterative design for the Library of Congress National Digital Library Program. *Proceedings of CHI '97*, March 1997, ACM New York. 518-525.

Rao, R. Pedersen, J, Hearst, M, Mackinlay, J., Card, S., Masinter, L., Halvorsen, P., Robertson, G., (1995). Rich interaction in the digital library, *Communications of the ACM*, 38(4), 29-39.

Shneiderman, B., Byrd, D., and Croft, B. (1997). Clarifying Search: A User-Interface Framework for Text Searches. *D-Lib Magazine, January 1997.* http://www.dlib.org/dlib/january97/retrieval/01shneiderman.html

Slaughter, Shneiderman and Marchionini, G. (1997). Comprehension and Object Recognition Capabilities for Presentations of Simultaneous Video Key Frame Surrogates. *Proceedings of the First European Conference on Research and Advanced Technology for Digital Libraries*, (Pisa, Italy, September 1-3, 1997), 41-54.

Smith, T. (1996). A digital Library for geographically referenced materials, *Computer*, May 1997, 54-60.

Wactlar, H., Kanade, T. Smith, M., Stevens, S., (1996) Intelligent access to digital video: Informedia project. *Computer*, May 1996, 46-52.