Previews and Overviews in Digital Libraries: Designing Surrogates to Support Visual Information Seeking

Stephan Greene*, Gary Marchionini[1], Catherine Plaisant, and Ben Shneiderman[2]

Human-Computer Interaction Laboratory,
University of Maryland Institute for Advanced Computer Studies
[1]also School of Information and Library Science, University of North Carolina at Chapel Hill
[2]also Department of Computer Science and Institute for Systems Research
College Park, Maryland 20742-3255, +1 301 405 2725
greene@las-inc.com, march@ils.unc.edu, {plaisant,ben}@cs.umd.edu,
http://www.cs.umd.edu/hcil

ABSTRACT

To aid designers of digital library interfaces, we present a framework for the design of information representations in terms of *previews* and *overviews*. Previews and overviews are graphic or textual representations of information abstracted from primary information objects. Previews act as surrogates for one or a few objects and overviews represent collections of objects. A design framework is elaborated in terms of the following three dimensions: (1) What information objects are available to users, (2) How information objects are related and displayed, and (3) How users can manipulate information objects. When utilized properly, previews and overviews allow users to rapidly discriminate objects of interest from those not of interest, and to more fully understand the scope and nature of digital libraries. This paper presents a definition of previews and overviews in context, provides design guidelines, and describes four example applications.

Keywords

Browsing, Information Seeking, Surrogates, Previews, Overviews, Visualization, Digital Libraries, User Interfaces.

INTRODUCTION

Moviegoers are enticed and informed by well-designed previews. Book-lovers roam the aisles of bookstores scanning jackets for titles and authors, and flip pages for catchy phrases or images. And now web surfers scan home pages or search result sets to decide where to go next. Unfortunately, in emerging digital libraries, designers often fail to provide appropriate views of materials to give an overall sense of the structure and materials available. This paper seeks to address these problems by recommending designs for *previews* and *overviews* as information representations.

We first define preview and overviews in context, propose three design dimensions, summarize design guidelines, then describe examples of interfaces making extensive use of previews and overviews, and conclude with recommendations for design and further research.

Interfaces for digital resources regularly fail to provide honest representation of what they include, wasting users' time and increasing their frustration with online systems. Some museum web sites give the impression of having large amounts of digitized art, but in fact only offers exhibit descriptions and links to other sites. Likewise, readers of corporate web sites often are led to believe that extensive product information and user support is available online, but following links may only lead to advertisements and a few frequently asked questions. Text summaries for Web search results remain too cryptic, indicating little of significance about the object previewed, nor its relation to the query as expressed.

Users prefer comprehensible, predictable and controllable environments in which they can rapidly and safely explore and use information. Design guidelines intended to help realize such systems abound, but recent work in the HCI literature has produced several characterizations of user interface design parameters that have advanced

* Current address: Language Analysis Systems, Inc. 2214 Rock Hill Road, Herndon, VA 20170 greene@las-inc.com

more thorough, concrete and practical concepts for the design space of interactive systems. Furnas (1997) proposed useful criteria for designing and evaluating representations of information structures in support of traversal and navigation. Ahlberg and Truve (1995) described the design space of query devices. Tweedie (1997) outlined several design dimensions for describing and comparing the interactive aspects of information representations. Shneiderman (1996) characterized information visualizations in terms of data types and user tasks.

Creation and evaluation of information surrogates are long-standing research issues in library and information science. Heilprin (1985) argued that the abstractions and compressions he termed *homomorphic reductions* are fundamental to information science; Borko & Bernier (1975) in their work on abstracting listed two dozen types of text surrogates. Surrogates are essential for automated and analytical searching because metadata such as Library of Congress MARC records, NASA DIF records, or data codebooks are often composed of standard sets of identifiers that searchers use to make decisions about what objects to examine next. In browsing, surrogates provide an important alternative to primary objects as they take far less time to examine and provide enough semantic cues to extract gist and allow users to assess the need for further processing of other surrogates and the primary object. In digital libraries and archives, surrogates are crucial for browsing large distributed collections that result from filtering programs or analytical queries of the data space. The need for inventing new types of surrogates underlies much of the research in digital libraries.

Our analysis of previews and overviews builds upon and complements these previous efforts, within the context of a theory of information seeking. Our analysis is based on the design and evaluation of several digital library projects (with NASA [Plaisant, Shneiderman & Doan, in press; Green et al., in press], the Library of Congress [Marchionini, Plaisant & Komlodi, in press; Nation et al., 1997; Plaisant, et al, 1997], and the National Library of Medicine [North, Shneiderman, & Plaisant, 1996]), and most of our examples are drawn from this work.

PREVIEWS AND OVERVIEWS IN CONTEXT

People seldom take the time to act optimally, often 'satisficing' rather than optimizing (Simon, 1979). Additionally, we make decisions without perfect information and according to experiential biases (Tversky & Kahneman 1974). Given its size and heterogeneity, information seeking needs in the WWW are almost impossible to perfectly satisfy, other than needs of simple fact-retrieval. Previews and overviews, if well designed, can aid users in making better decisions about potential relevance, and extract gist more accurately and rapidly than traditional hit lists provided by search engines.

Previews and overviews are graphic or textual representations of objects of interest. As a first approximation:

- A preview is extracted from, and acts as a surrogate for, a single object of interest.
- An overview is constructed from, and represents, a collection of objects of interest.

The design of previews and overviews is guided by the view that information seeking is a dynamic and iterative decision-making process in which users engage to change their knowledge state (Marchionini, 1995). This view holds that except for known-item searches, information seekers initiate search through a query or by browsing in a subset of the digital library. Users scan objects rapidly to determine whether to examine them more closely, or move on in the digital library. This process continues until the information need is satisfied or satisficed or the search is abandoned.

This view of information seeking requires system designs that provide multiple levels of representation for information objects. The levels may range from terse but informative surrogates such as titles, to increasingly detailed surrogates that become successively more indicative of primary objects. To help ensure the success of such processes of progressive refinement, it is important to provide users with appropriate intermediate surrogates that facilitate the rapid elimination of objects from further examination. Our design strategy provides users with overviews of collections which are used to select successively more detailed overviews or specific previews. Where appropriate, an overview can also provide direct access to primary objects.

Previews are analogous to bibliographic records and overviews are analogous to catalogs. It is important to note that a catalog may itself have a bibliographic record (preview) and a bibliographic record may serve as a catalog (overview) for a complex object such as chapters of an anthology. Depending on the user’s task and the nature of the information represented, some designs fit the definitions of both previews and overviews, or fit one definition and then the other at different points in the interaction. For example, the thumbnail image of the Visible Human Explorer (top right of Figure 1) acts as a preview when the user is selecting images to download, but might be used as an overview to browse the full size image in a classic overview/detail image browser (Plaisant, Carr, &

Shneiderman, 1994). Any information surrogate will embody, to some degree, preview or overview aspects. The preview-overview

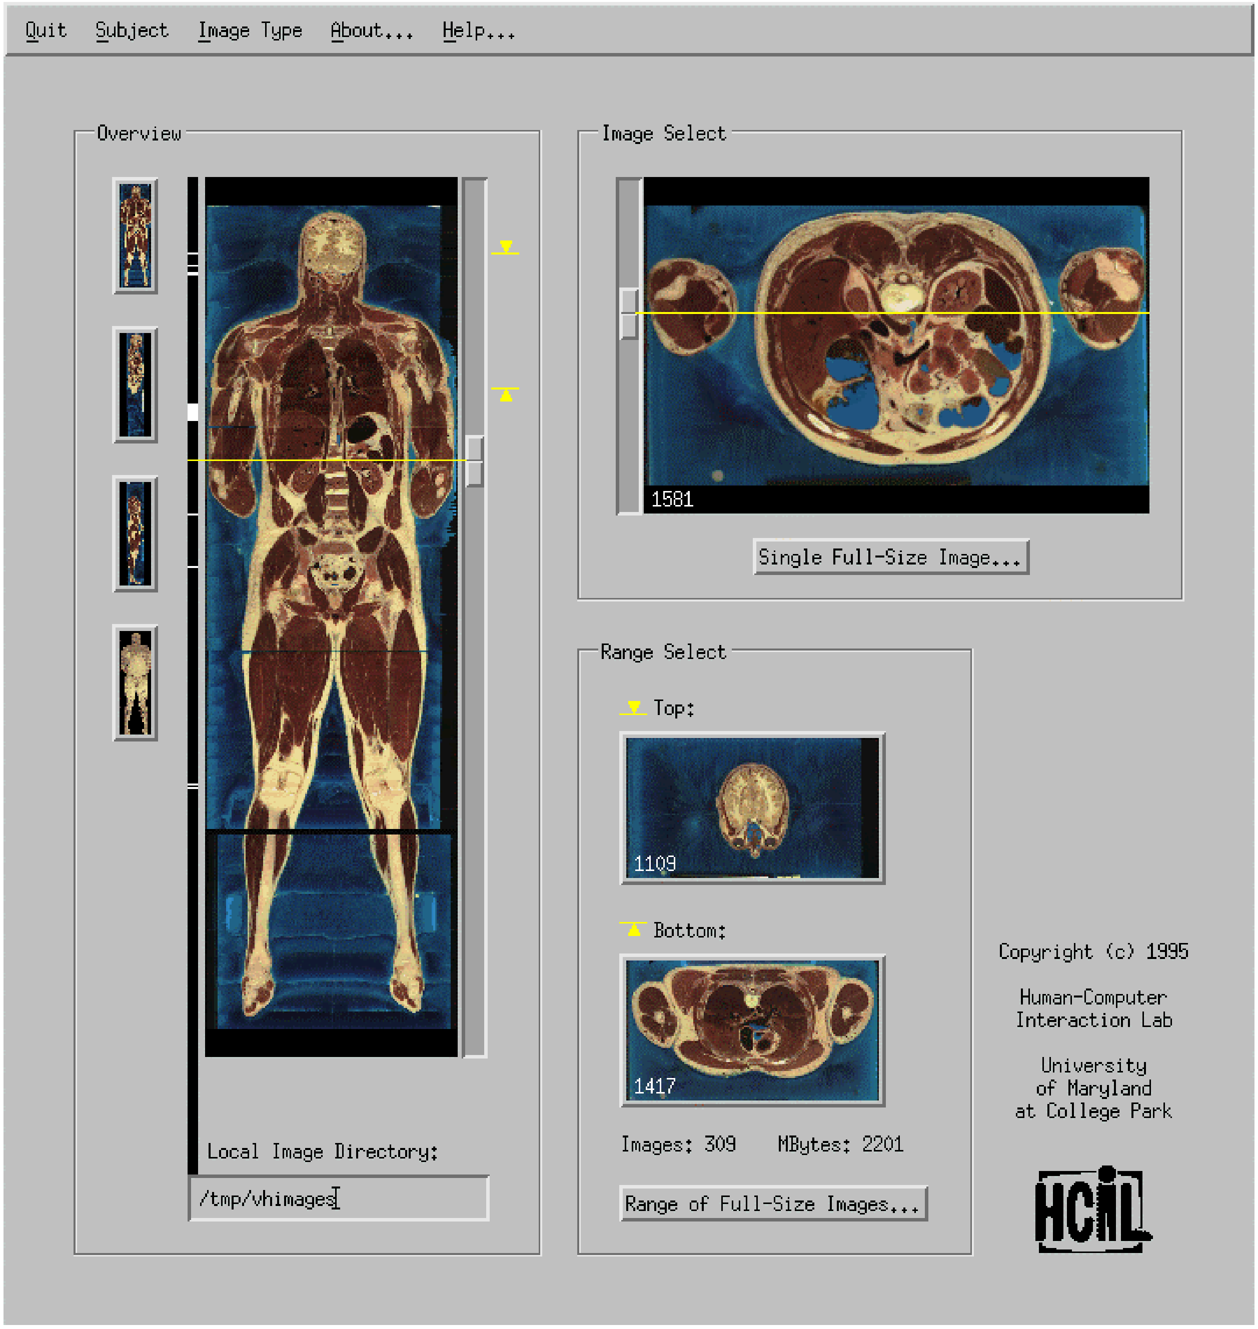

Figure 1. The Visible Human Explorer (North et al., 1996) allows users to rapidly browse thumbnail previews of high-resolution images from a digital library that would take days to download and review. On the left, a composite full-body thumbnail overview, constructed from the 1800 separate images, allows users to browse individual image previews in a matter of seconds.

characterization thus provides a fundamental dimension along which to design and evaluate a wide variety of information representations.

Browsing is a human activity that closely couples search and use as parallel activities. Unlike analytical search strategies, which can be batched and automated, browsing is an inherently interactive process that demands ongoing human attention. Thus browsing is ultimately constrained by the limits of human physiology and cognition. Well-designed surrogates help exploit human perceptual strengths while not exceeding cognitive limitations. Human processing will be aided by surrogates that leverage natural mappings between data and representation (e.g., spatial data on maps rather than in tables, interval data on properly proportioned scales, etc.)

Previews and overviews provide a useful family of information representations that can help to meet these challenges. An effective preview is an information surrogate that communicates to the user, at the appropriate time, sufficient information about the primary object it represents to support users in making a correct judgment about the relevance of that object to the user's information need. An effective overview provides users with an immediate appreciation for the size and extent of the collection of objects the overview represents, how objects in the collection relate to each other, and, importantly, what kinds of objects are not in the collection. Well-designed previews and overviews will also provide users with directly manipulable control mechanisms that facilitate exploration and allow users to focus attention on the larger information problem at hand.

DESIGN DIMENSIONS OF PREVIEWS AND OVERVIEWS

The design of effective previews and overviews requires a thorough analysis of the application domain with respect to three interdependent dimensions:

- What information objects are available to users
- How information objects are related and displayed
- How users can manipulate information objects

These dimensions guide the design of surrogates by helping to determine what is to be represented, how representations are to be organized and displayed, and how the representations are controlled during interaction. However, this does not imply that surrogate representations are organized and manipulated in ways analogous to the way primary objects are organized and manipulated. Especially for overviews, innovative representations will be abstractions of primary objects and their attributes, and thus often require their own methods of organization of manipulation. Nonetheless, effective designs result from the careful analysis of primary objects and users' tasks. Design guidelines for previews and overviews emerged from our experiences in building interfaces that applied user-centered HCI principles and practices. Successful designs were achieved with methods that were both top-down, based on general design principles, and bottom up, using good sense and iterative testing.

What Information Objects are Available to Users

The optimal identification of information objects, and the design of appropriate representations for them, requires carefully balancing of the demands posed by data characteristics, user needs, and system constraints.

Identifying Primary Objects

Primary information objects should be identified from a user-centric perspective. They are logical entities useful to users, not necessarily identical to entities as stored in the system. For example, although an image and its caption may be physically retrieved from separate system files, they are likely to be considered as a single object by users.

When primary objects are too diverse or complex, an aggregate form of those objects can be chosen for use in previews and overviews. In our work with NASA (detailed below), the "ultimate" primary object useful to the user is observation-level earth science data. But the data comes in such great variety of forms (numeric observations, images, series, etc.) that objects cannot be uniformly described. Dataset records were a better choice of primary object for the initial query interface because they have consistent metadata that can be used to create surrogates (Figure 2a and 2b).

The user task determines the most appropriate level at which to consider an information object. That is, function establishes the optimal granularity of the relevant data type. In the NASA example, scientists need to identify data to conduct a particular inquiry. They are better served by an interface that lets them start by reviewing and selecting a small number of datasets, as opposed to wading through millions of objects. For the Library of Congress National Digital Library Program, the diverse user community suggested that we provide multiple entry points.

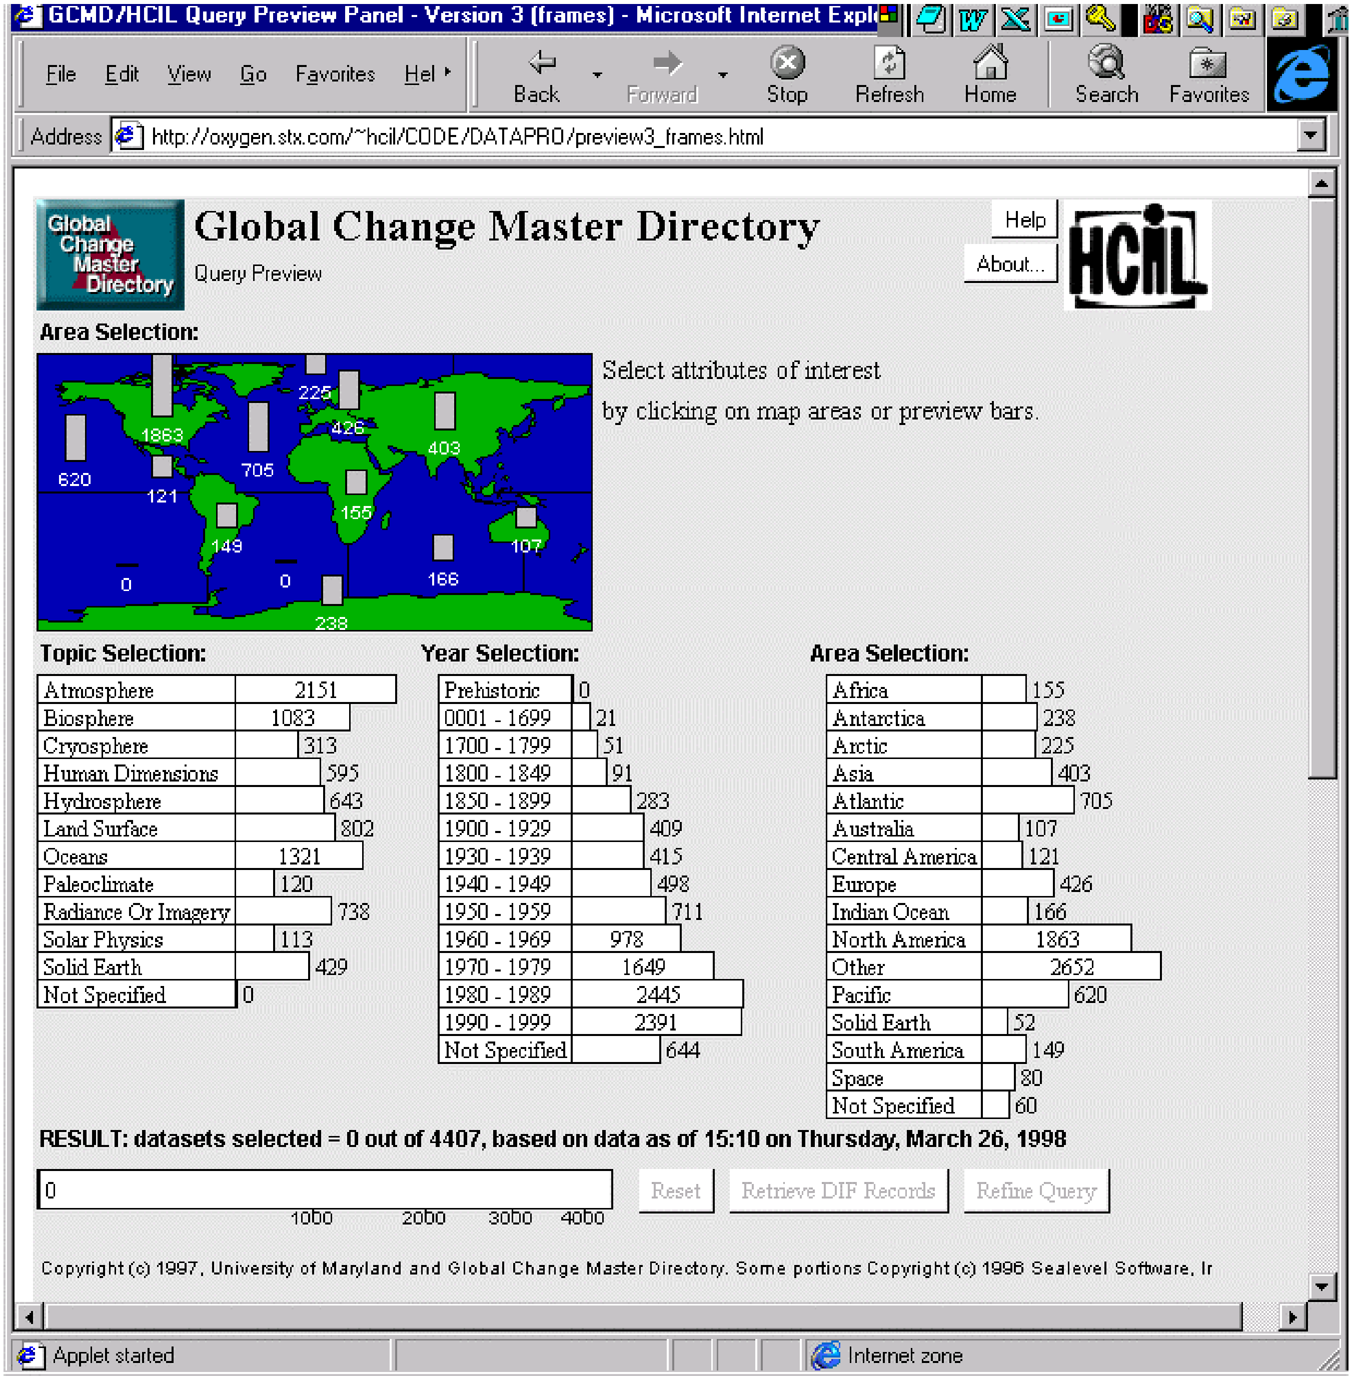

Figure 1a. The interface for NASA earth science data provides a single screen overview of the digital library (i.e. here the numbers of datasets available in the NASA Global Change Master Directory.)

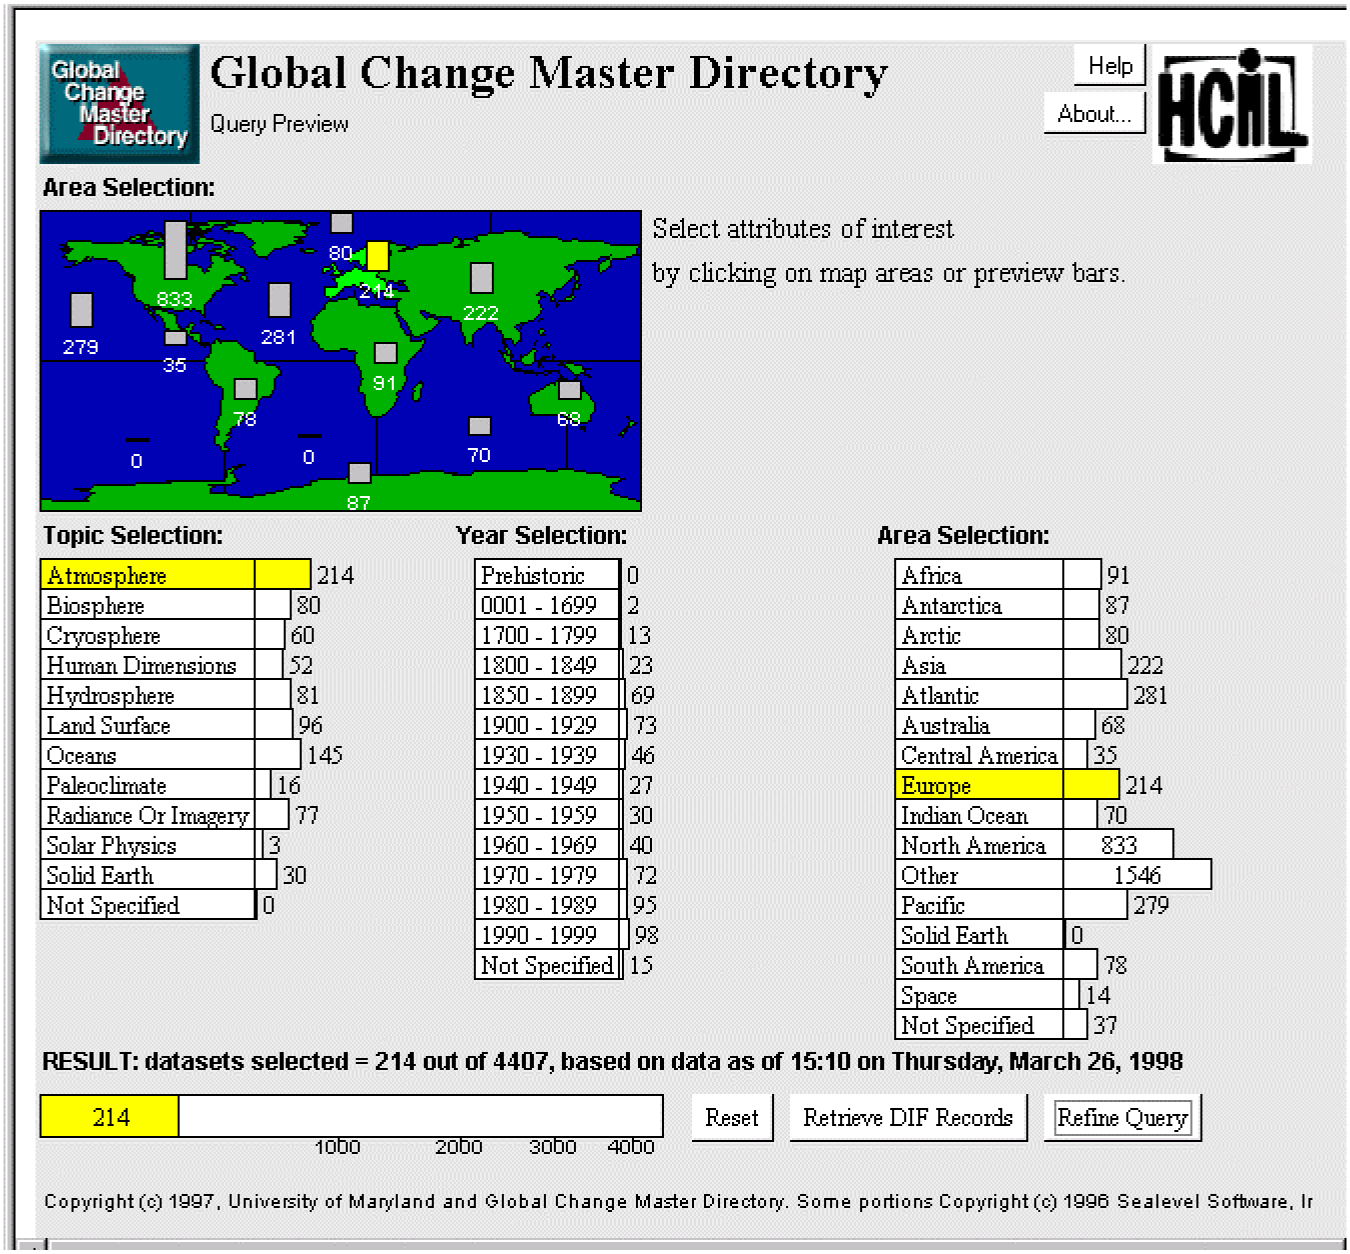

Figure 2b. When users select attribute values (here atmosphere and Europe), count totals are updated using precomputed count preview tables. This gives users a preview of the size and distribution of the result set before the query is submitted, and eliminates the zero-hit or mega-hit query problem.

Defining the Role of Object Surrogates

Surrogates for objects are created to serve several purposes. A primary motivation generally dominates the design process, although additional purposes are usually served in tandem. Many surrogates may already be available to the user interface designer because they are inherent to the data (e.g. title or section headings in structured documents) or because they were created to serve other system requirements (e.g. bibliographic catalog entries are created for information management and retrieval purposes and are also useful as basic previews). In all cases, an information surrogate adds value that complements the primary information object.

The main roles for preview and overview surrogates include:

- Aiding retrieval (i.e., used as the basis for indexes)

- Aiding users in quickly making relevance decisions
- Reducing network data transfer (i.e. reducing time and resource needs) to make practical the examination of numerous resource-intensive primary objects
- Supporting the user's need to capture the gist of complex heterogeneous information
- Providing indicators of scope, size and structure of large information spaces
- Informing and enticing with representative samples, while helping users define productive queries

Options in the Surrogate Design Space

Data perspective: Surrogates may be literal extracts which are extracted without change from the primary object, or original constructs, which are alterations or original creations that convey meaning from the primary object. This distinction is similar to the distinction between constructed and converted data as defined by Tweedie (1997). Examples of literal extracts include thumbnails, tables of contents, pull quotes, video streamers (Elliot, 1993), low-resolution images, and sampled video frames. Examples of original constructs include text abstracts, controlled vocabulary term lists, collage images or salient stills (Teodosio & Bender, 1993). Original constructs are generally more difficult to create automatically than literal extracts, and thus are often considered to have "more" added value than literal extracts.

This dimension of surrogate design shows that the original definition of previews as extracted surrogates and overviews as constructed surrogates is relative to the degree of content transformation (opposed to system operation or level of effort). For example a randomly generated collage of sample images from a photo collection may be constructed but does not require any intelligent intervention and remains a literal surrogate.

User perspective: From the user’s point of view, previews are very sensitive to time, while overviews are less so. Users preview objects before examining their details, and usually need not reference the preview again. In contrast, users can make good use of overviews before, during, or after examining details.

Preview and overview surrogates can give users representations that differ in their degree of abstractness and exhaustiveness. Abstractness refers to the degree to which a surrogate is removed from a literal representation of the primary object. The vivid images of the Visible Human Explorer (Figure 1) sit toward the highly literal end of the scale and the bar graphs of the NASA Query Preview interface (Figures 2a and 2b) toward the other end. Exhaustiveness is a related concept that refers to the degree to which a surrogate reflects the entirety of the primary object(s) it represents. Overviews by definition should be fairly exhaustive, aiming to represent the most salient features of objects in the collection, but problems like overlaps or a limited number of represented attributes might limit the coverage of the overview. Previews are much less likely to be exhaustive.

Additionally, previews can influence the attention of users as they decide whether to look at the primary objects. A preview has the potential to be misleading, eliciting imprecise or incorrect generalizations about the primary object. Overviews, by contrast, can be used to understand an object in context. It is safer to generalize about an object from an overview because it aims to be a general summary or description, and supports comparison among related objects. Thus, in general, previews may be more affective while overviews tend to be more cognitive.

How Information Objects are Related and Displayed

Successful browsing depends on the attributes of individual units of information, and on how different units are related to one another and displayed to users. Effective previews and overviews take advantage of data type-specific structures. Preview designs are primarily constrained by what attributes are available, and design decisions center on selecting attributes from among those available, and determining how best to represent them in a surrogate. Overviews are more challenging as their design depends on commonality of attributes across objects. Design decisions are often complicated by a scarcity of common attributes and the need to create artificial summary attributes.

Browsing previews is facilitated by regular object structures such as single or multiple columns, linear sequences of scenes or MTV-like montages, and comma delimited or tabular layouts, which can be used in the design of the surrogate. A good example is TableLens (Rao & Card, 1994), which uses a graphical overview that mirrors the original information object (i.e. the numeric table).

Browsing overviews of collections of objects with common attributes (e.g., a list of retrieved records from a database or online catalog) allows users to exploit consistent surrogate structure displayed in lists, tables, and other common layouts. However, search engines and filters that work across many information sources often yield heterogeneous objects. A result set may include representations of primary objects such as full texts, videos, or datasets employing different levels of object surrogates such as titles, abstracts, samples, or summary statistics. For example, on the Web, users may have to browse vertically across several levels (e.g., across specific pages, sites, or indexes of sites). Another example is a query in the Library of Congress National Digital Library that may return a set of hits that mix collections as a whole (e.g., the name of a photograph collection), finding aids (e.g., a full-text document that describes and gives access to a manuscript collection), and individual objects (e.g., a specific photograph or sound recording). In these cases, common attributes may be absent and more abstract attributes must be formulated to facilitate overviews.

In all cases, the goal is to create previews and overviews that communicate the level and scope of objects to users so that comprehension is maximized and disorientation is minimized. A growing number of surrogates address this problem, such as in our Library of Congress National Digital Library prototype (Marchionini et al., in press; Plaisant et al., 1997).

How Users Can Manipulate Information Objects

The third major design dimension is user control of surrogates. We advocate mechanisms that put users in control without interfering with their progress. Direct manipulation and information visualizations help enormously in this regard. The manipulation of individual surrogates can be considered in terms of interactivity and tailorability. For example, the interactivity of a preview of a collection consisting of a collage of samples can range from a simple static graphic, to an interactive slideshow of samples, each navigable to its corresponding primary object. Tailorability of surrogates can be vital in complex applications that employ large numbers of attributes for complex information objects. We employed overviews that allow users to choose which attributes dominate the surrogate representation, or to modify display attributes such as color, size or labels. Given a rich set of interactivity and tailorability mechanisms, an overview surrogate can become a complete application, such as a dynamic query interface coupling a zoomable overview with a real-time filtering mechanism using sliders and buttons (Ahlberg & Shneiderman, 1995).

In the dynamic and interactive process of information seeking, appropriate control mechanisms are also needed to navigate *among* surrogates at the same or different levels of representation. The two classic control mechanisms are scroll and jump but zooming is becoming a viable alternative mechanism as machines become more powerful (e.g., Bederson & Hollan, 1994). Zooming offers users the continuity of scrolling as well as a natural metaphor for browsing vertically across different levels of surrogates.

Another technique of control is the use of multiple coordinated windows or views for presenting surrogates. Current work enhances information coordination and organization across multiple windows (North & Shneiderman, 1996). The Visible Human Explorer is a good example of an overview coordinated with a preview (Figure 2). Users can select a preview by sliding the cursor on the overview, or change the overview forward and backward by sliding the cursor of the preview.

Ever more specialized manipulation mechanisms are being developed for certain data types. For example, video previews (Wactlar et al., 1996) can be presented one or several at a time, at a speed adjustable by users (Slaughter, Shneiderman, & Marchionini, 1997). Empirical data on optimal frame rates for video previews can guide designers of video resources (Ding, Marchionini & Tse 1997). Additional empirical data on user performance with alternative representations and control mechanisms for temporal information is needed.

DESIGN GUIDELINES

The following design guidelines grew out of our experience in building and testing interfaces. The guidelines are helpful in applying the concepts of previews and overviews to novel contexts defined by particular users, data, and systems.

- Use salient surrogates. This requires a solid understanding of data and users (task analysis and needs analysis).
- Use multiple surrogates. This supports different user styles and experiences.
- Use multiple levels of surrogates. Display them on a single screen, where possible, in order to show hierarchy and minimize mouse clicks.

- Use surrogates to inform users about size, extent, and availability of collections or objects.
- Leverage data types (use visual surrogates for visual data, etc.).
- Choose metaphors that map onto the data or most common task.
- Provide logically sequenced surrogate levels and capitalize on natural orderings (e.g., geography, chronology).
- Provide directly manipulable control mechanisms (visual, incremental, and reversible controls) for navigating both within and across surrogates.

EXAMPLES FROM APPLICATIONS

In addition to the Visible Human Explorer (Figure 1) we present three case studies of interfaces that illustrate the application of the design guidelines to diverse interfaces employing previews and overviews. We present in more details the NASA query preview example (Figure 2a and 2b), then describe our work for the Library of Congress National Digital Library Program, and the LifeLines interface designed in the context of juvenile justice and medical record applications.

NASA EOSDIS Query Preview Prototypes

NASA's Earth Observing System Data Information System (EOSDIS) allows users to retrieve earth science data from hundreds of thousands of data resources containing images, measurements, or processed data from data archive centers. Metadata is available and is used to search the archives. Standard EOSDIS metadata includes spatial coverage, time coverage, type of data, sensor type, etc. Traditional form fill-in interfaces for EOSDIS permit searches of the large holdings but zero-hit and mega-hit queries are a problem. When results are returned, it is difficult to estimate how much data is available on a given topic and what to do to increase or reduce result set sizes. Our approach to these challenges is to employ overview and preview surrogates of highly abstracted metadata that allows users to rapidly and dynamically eliminate undesired datasets. The reduced volume of the abstracted metadata allows queries to be previewed and refined locally before they are submitted over the network.

Our query preview interface (Figure 2a) features a single-screen overview of all earth science data available from a designated information service provider, the Global Change Master Directory (GCMD)[Plaisant et al, in press; Greene et al., in press]. The data are characterized with three high-level attributes: location, years of coverage, and general topic area. These three attributes were generally agreed to be the most salient attributes associated with earth science data resources. For each of the three attributes, the large set of attribute values was aggregated up to around ten high-level attribute designations. The granularity of the attribute values chosen for use in the interface is deliberately crude, in order to be able to represent vast amounts of scientific data in a single overview. The careful choice of appropriate attribute and attribute value aggregations represents a significant portion of the design effort and reflects the needs of scientific users. Following the guidelines we used multiple surrogates (three) that leveraged data types (geography and chronology) to provide a rich overview of the digital library. Visual surrogates were used when possible (i.e. a map, and bars of varying lengths for the counts).

For each attribute value in the interface, the number of available data products bearing that attribute value is shown. This gives users an overview indicating the size and extent of the available data. By clicking on combinations of attribute values, users can prune to reflect only those data items that satisfy selected attributes. In this way, users generate a preview of the result set that would be returned if a query matching the selected attributes were sent over the network[1]. Zero hit queries are thus eliminated and users can explore the data without penalty or network delay. Following our guidelines, the controls were made visual, incremental and reversible.

When the preview indicates a reasonable result set size, users submit the query with a mouse click rather than typing, which returns the list of dataset records, or move the query to a refinement phase. In the current operational version at GCMD, this consists of an additional overview phase that operates on the subset of data items from the first phase. Thus, what was used as a preview of a result set from the entire archive is now

---

[1] Accurate joint distributions for the attribute values are calculated from the local metadata using specialized data structures and algorithms, although these methods will not scale beyond the small size of the high-level database currently in use. Ongoing research is devising algorithms and data structures for large-scale application of these techniques [23].

represented as an overview of a new subset of the digital library. At this point, users are able to discriminate data objects using three additional attributes. This progressive refinement, where overviews lead to previews of subsets which in turn serve as overviews for more fine-grained objects illustrates how interactive overviews and previews can scale to large, complex digital libraries. An empirical study of a query preview interface demonstrated that user performance was improved for broad search tasks (Tanin et al., 1998).

While Furnas (1997) discusses methods for achieving effective view navigation, the NASA query preview prototype meets Furnas's criteria within the analogous context of data browsing. Furnas states that high-level semantics play a dominant role in designing navigable structures, and the success of the query preview prototype stems from the capture of attribute metadata at a carefully chosen, highly abstracted level. In addition, Furnas outlines the need to carve up an information space in a way that supports efficient navigation among distinct objects. Overlap patterns of "to-sets," the set of desired objects that may be reached by navigating across a link, must be able to "whittle down the space of alternatives in a small number of intersections." If one considers the records as the primary objects of interest in the query preview prototype, then the selection of attributes is a way of navigating toward the dataset records that fit the user's need. In the overview state, before selection occurs, all datasets are represented in the count totals for each attribute value. In the preview state, users are shown the degree to which they have moved closer to their desired result by the revised count totals. The attributes as chosen carve up the information space, and user selections quickly whittle down the space to a small set of objects.

Furnas also discusses the value of *residue*, referring to the degree to which a node in a navigable structure has the property of indicating what nodes are linked to it. The query preview interface has residue from all the primary objects it represents, in the count totals. In addition, the high-level attribute value bars indicate the nature and distribution of the properties of the represented objects. From this perspective, the design of preview and overview representations is essentially the task of determining what residue is to be captured, and how.

Library of Congress National Digital Library Prototypes

We have worked with a team at the Library of Congress (LC) to develop and test interface designs for LC's National Digital Library Program (NDLP) [Marchionini et al, in press; Plaisant et al., 1997). Our effort focused on the American Memory collections, which consist of historical multimedia materials: photographs, voice recordings, catalogs, manuscripts, journals, videos, etc. Five million objects are expected to be digitized by 2000. Some will be richly described with metadata such as bibliographical records, but the volume of materials is so great that a large proportion of the individual objects will not be described at all (and therefore never "hit" by traditional search). Collections of objects are more likely to be described and browsing of the undescribed objects of those collections will be important, just as it is in physical collections at LC such as those containing photographs.

To relate and display these objects our designs emphasize providing previews and overviews of materials to users. We developed several novel browsers that address the problems of navigating the large and diverse LC collections. For the complete family of American Memory digital collections, we developed a Collections Browser (Figure 3) applying dynamic query principles to the access of collections. For specific collections within American Memory, we developed a generic collection browser (called WebTOC, Figure 4), especially designed to be usable for both unprocessed and processed collections. Finally we designers specialized browsers to examine homogenous sets of objects such as images, pages or video. These browsers provide overviews of sets of collections or objects, and help users learn about topical coverage and the kinds of available materials. Users form concrete and realizable expectations, and can then formulate successful search queries or navigate with confidence.

Collection browser: The Collection Browser (Figure 3) uses dynamic queries (Ahlberg & Shneiderman, 1994) as its basis for user manipulation of objects. In agreement with our guidelines the controls are incremental and reversible, and multiple surrogates are used. The interface features a visual overview of all the collections: a zoomable timeline with a bar for each collection showing its time coverage, therefore leveraging the temporal data type. A set of widgets (three menus for location, topics and format attributes) provides a visual representation of query attributes. All surrogates are tightly coupled (i.e. timeline, list of collections and short descriptions are updated simultaneously). Users browse the digital library by interacting with widgets (setting attribute filters, zooming the timeline, or simple cursor selection). Each change is in effect a new query, and the resulting collections are immediately and continuously shown.

Designing a useful overview requires that some consistent metadata exist for all the primary objects of the overview. Early design work found that consistent metadata was not available for most materials at the object level, making it impossible to apply the overview design at the object level. Collections were therefore chosen as primary objects. But even there the metadata was inconsistent across collections. It became clear that it was important to define a set of consistent attributes at the high level of the collection if any effective browsing across the collections was to be provided. A metadata format was defined by the LC team, and these data were generated for the collections used. The timeline, collection list and menus are created automatically from this metadata. The design of this new surrogate level—an overview of the NDLP that allowed users to explore collections--was a significant part of the design process.

WebTOC: The individual collection browser tool we developed is WebTOC (Figure 4) [Nation et al., 1997]. WebTOC generates a table of contents using two different strategies: following existing HTML links or using the underlying file system structure. Following links is appropriate for existing collections, while using the file system is appropriate for newly digitized collections, which have not yet been linked, indexed or annotated (i.e. a physical box of items was merely digitized and saved in a subdirectory.)

A two-frame layout allows users to expand or contract the TOC presented in outline on the left frame, while inspecting any selected object in the right frame. As summarized in the WebTOC legend (top left of Figure 4), individual objects are represented with individual lines while bars aggregate the files behind a link or directory to represent the total size of the included objects. Color is used to represent file type, length, the overall size, and the shadow below the bar is proportional to the number of objects included. The number of documents and their size is displayed in the browser status area when the user's cursor is over the bar. Figure 5's TOC shows that some parts of the site are very small (nearly empty) relative to others. It would have been impossible to learn this from the home page on the right. In addition, a small shadow suggests that only a few objects exist under a topic. WebTOC can also be used to display results of a search in the context of the structure of the collection, opening only branches of the tree where objects where found.

Following our guidelines, WebTOC uses multiple surrogates, made visual whenever possible (type was mapped to color, and size to bar length, thickness or shadow size.) Several logical levels show objects within their parent containers (subdirectories or the entire collection). In this example, there are very few usable attributes because of the missing metadata at the object levels--a common challenge for digital libraries.

We designed and refined WebTOC to satisfy our own need to understand the huge directory of the Coolidge collection—what it contained and how it was organized. WebTOC proved to be useful to us in working with the unprocessed Coolidge collection and to library staff in preparing or reviewing new collections. It is likely to be useful for the general public as it provides an alternative paradigm for access to the materials. It is particularly helpful for exploring a collection when searching is inadequate due to the lack of metadata, finding aids, or searchable text; or because users are not able to specify their search appropriately. It allows users to discriminate at a glance large from small collections, or collections made primarily of links from sites including digital materials. The graphical overview gives an appreciation of the type of materials. Finally, 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, and can be used to display results in the context of the hierarchy. WebTOC is an exhaustive overview tool that can help users to browse, and library staff to develop and organize the collections.

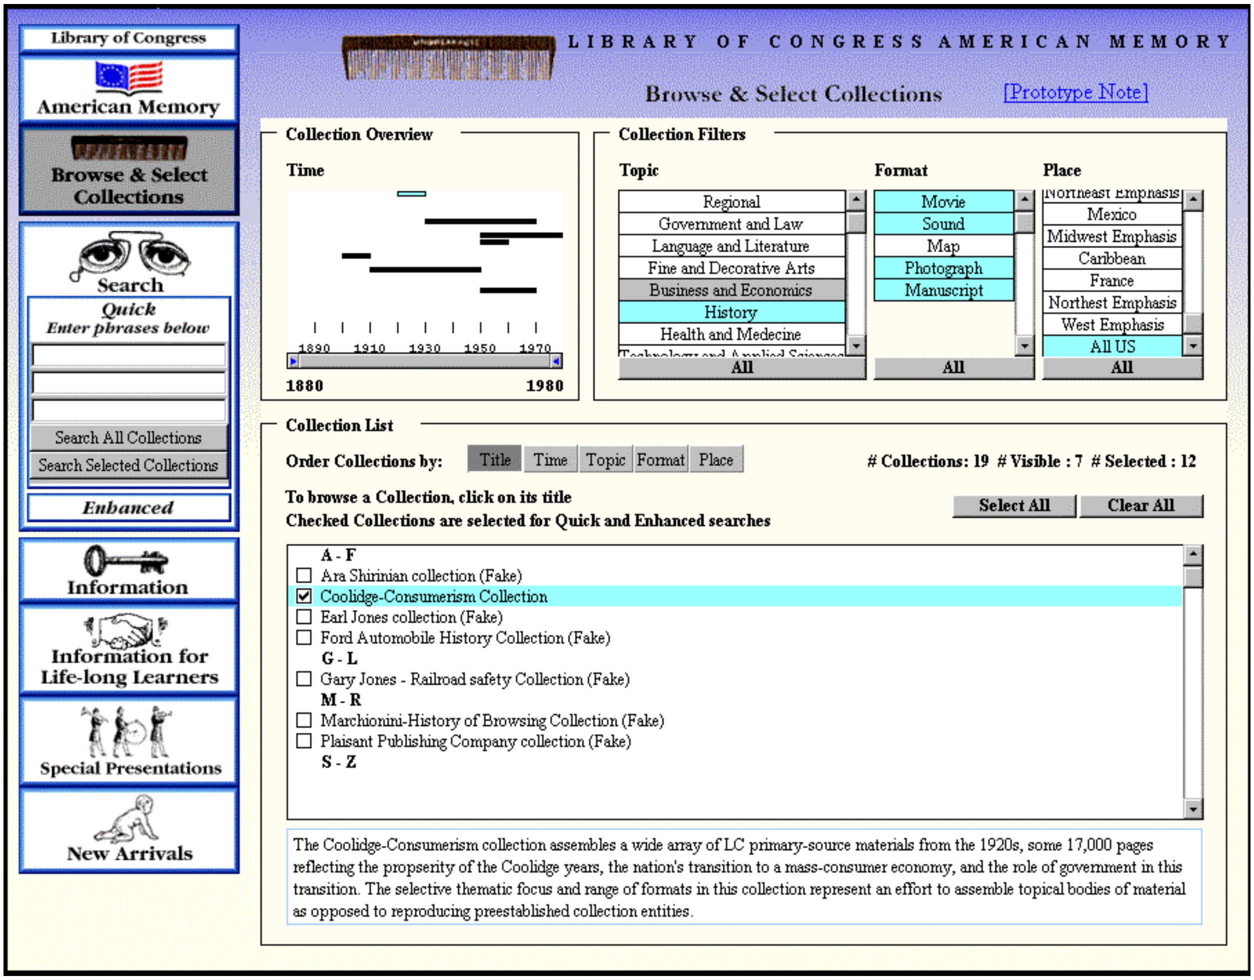

Figure 3: The NDLP Collections Browser. The overview showing the time coverage of the collections has been interactively filtered to show only collections including manuscripts on the topic of business and economics. Multiple surrogates are dynamically controlled with visible, incremental and reversible controls. When users mouse over a collection's title or timeline bar, all other surrogates are updated as well i.e., the short description is updated, and the title, timeline bar and attributes are highlighted.

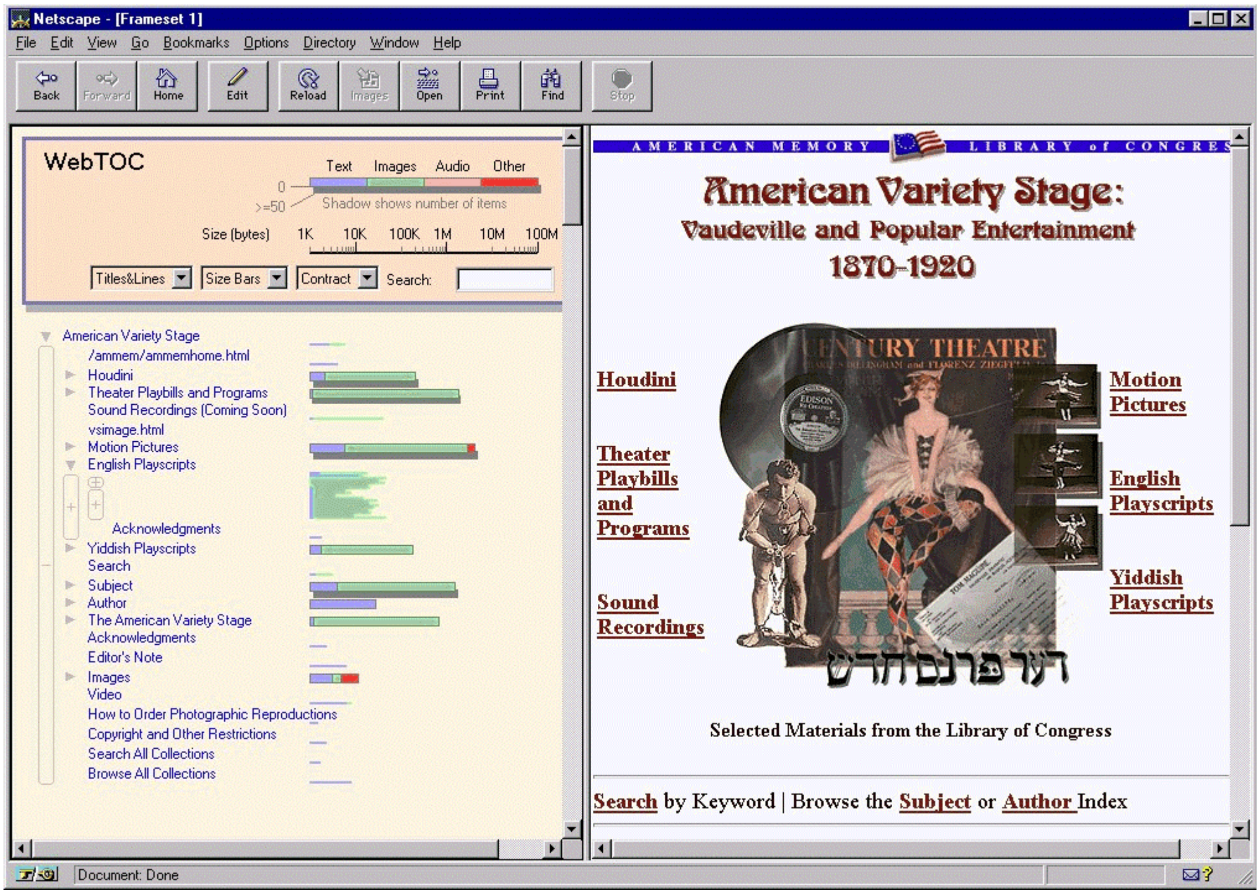

Figure 4. The WebTOC overview of the American Variety Stage Collection. The outline can be expanded to reveal more details. Multiple surrogates are displayed using visual and textual characteristics. A control panel allows user to customize the display.

Search results browsers: We designed a series of preview tools allowing sets of homogenous materials to be reviewed at once.

- An image browser lets users preview thumbnails of search results or collections (Figure 5).
- A video browser allows users to examine sets of videos, each one previewable at the user's option as a set of animated key frames or storyboard of keyframes.
- A page-oriented browser allows users to peruse journals, manuscripts, and other linear sets of page images. On a single screen, surrogates such as thumbnails, table of contents, finding aids, etc. as well as full views of the pages can be browsed.

These examples illustrate that when consistent information objects of a given type are available, they can be related and displayed in abundance in just a few screens, and simple control mechanisms afford rapid perusal. When results are heterogeneous and metadata inconsistent, browsers remain simple lists of objects. In our prototype, heterogeneous lists were summarized by type (e.g., 25 images, 5 videos, 7 journals etc.) and links to specialized browsers provided for those result subsets. In addition, next to each object a list of all the available surrogates for the object was given.

The American Memory digital collections contain enormous volumes of material in many different formats and with different levels of descriptive metadata. This makes query-form-based searching difficult and browsing more important. Representations of these diverse information objects were made available to users in a variety of forms. A preview of a collection can be samples from that collection (Figure 5), while a preview of a photograph can be a thumbnail. An overview of all the collections displays on one screen each of the collections with some of their common attributes. An overview of all the objects in a collection (webTOC) shows a representation of all objects, with all or part of the available common attributes made visual. These examples illustrate the end results of user needs assessment, task analysis, and ongoing interactions with LC staff. They demonstrate various ways that our general design guidelines for previews and overviews may be instantiated in practical digital libraries.

Lifelines

Another example of the benefits of overviews can be see in the LifeLines interface. LifeLines (Figure 6) is a visualization interface for collections of personal history records. It was used in the context of juvenile Justice for delinquent youth records (Plaisant et al., 1995) and to medical records (Plaisant et al., 1998). Timelines provide an exhaustive overview and direct access to the details as the overview acts as a giant menu. Figure 6 shows the example of a medical record. The display is organized in groups or facets (problems, allergies, diagnosis etc.) Lines correspond to continuing events (e.g. showing that the patient quit smoking in January). Line thickness can indicate severity, color can indicate status (e.g. a abnormal test result). Test, procedures or drugs linked to a diagnosis can be highlighted. For a juvenile justice youth record, cases, placements or assignments can be represented as horizontal lines, while icons represent discrete events such as reviews or hearings.

LifeLines was designed to provide an exhaustive overview and act as a giant menu as each element of the display leads to the corresponding detail document(s). The primary information objects are events, which have various attributes, including their start and end times. The simple metaphor of timelines clearly relates events to each other. Rich displays and interactivity allows users to evaluate attributes via color, size, label coding, or filtering and zooming. Rich control panels and zooming of the time scale provides for tailoring of the display according to users' need (Plaisant, Shneiderman & Mushlin, in press). The data model includes an aggregation mechanism that allows sets of events to be summarized. The summary events acts as surrogates that are progressively transformed when users zoom, until the underlying primary objects appear. For example, a simple line for a hospitalization will become two lines for the emergency room and the regular admission, each finally revealing the details of the tests and procedures.

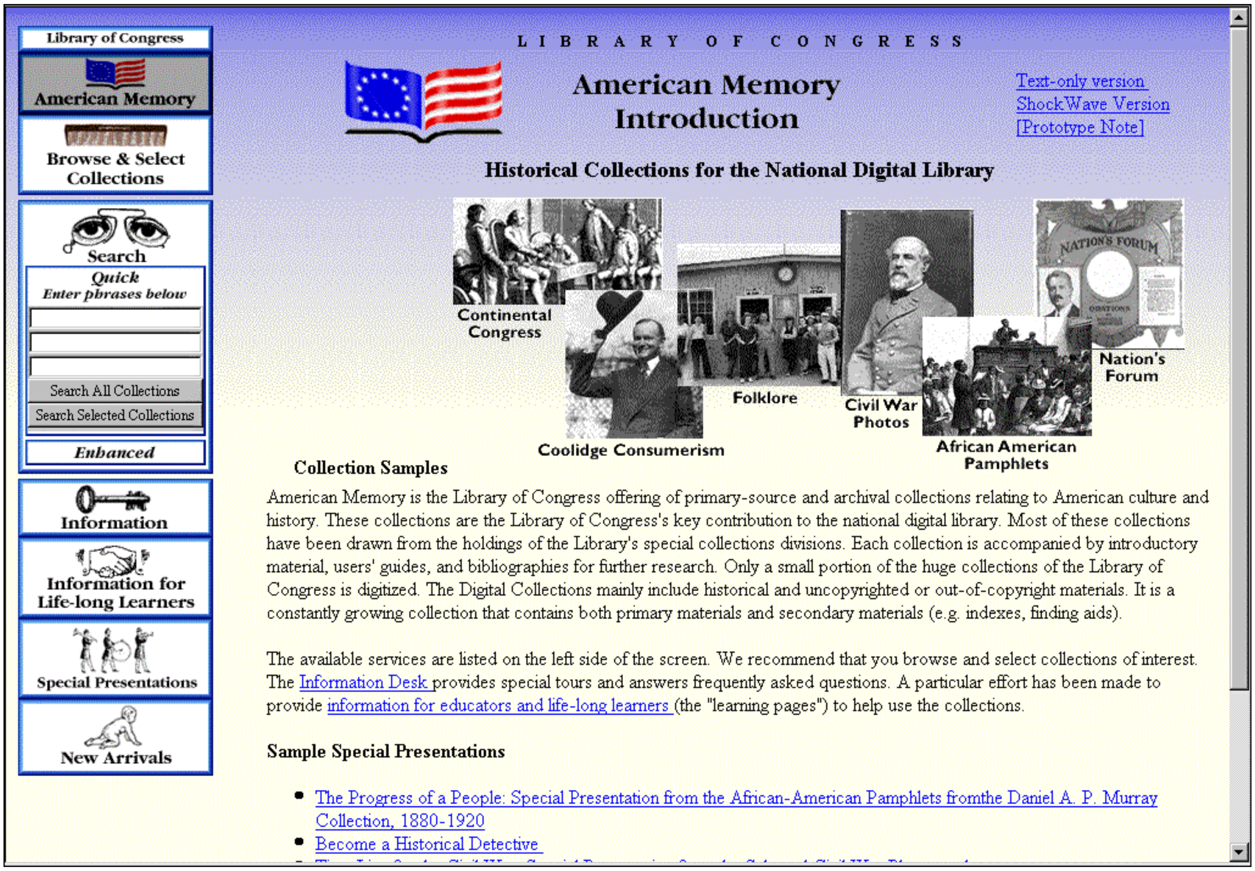

Figure 6: Opening screen for the American Memory Collections prototype, showing animated samples of the collections. Samples could be carefully chosen by a librarian to highlight hot or most representative objects, or selected automatically at random or among the most often requested objects. Animation allows a more exhaustive preview to be presented and is visually attractive. Samples bring treasures to the surface and provide an appreciation for the diversity of materials and the topics covered.

Figure 5: This simple image browser shows a series of previews. Showing 50 thumbnails at once allows users to browse a 200-photo collection in 4 screens. Images can be selected and viewed in larger size, or saved in personal collections.

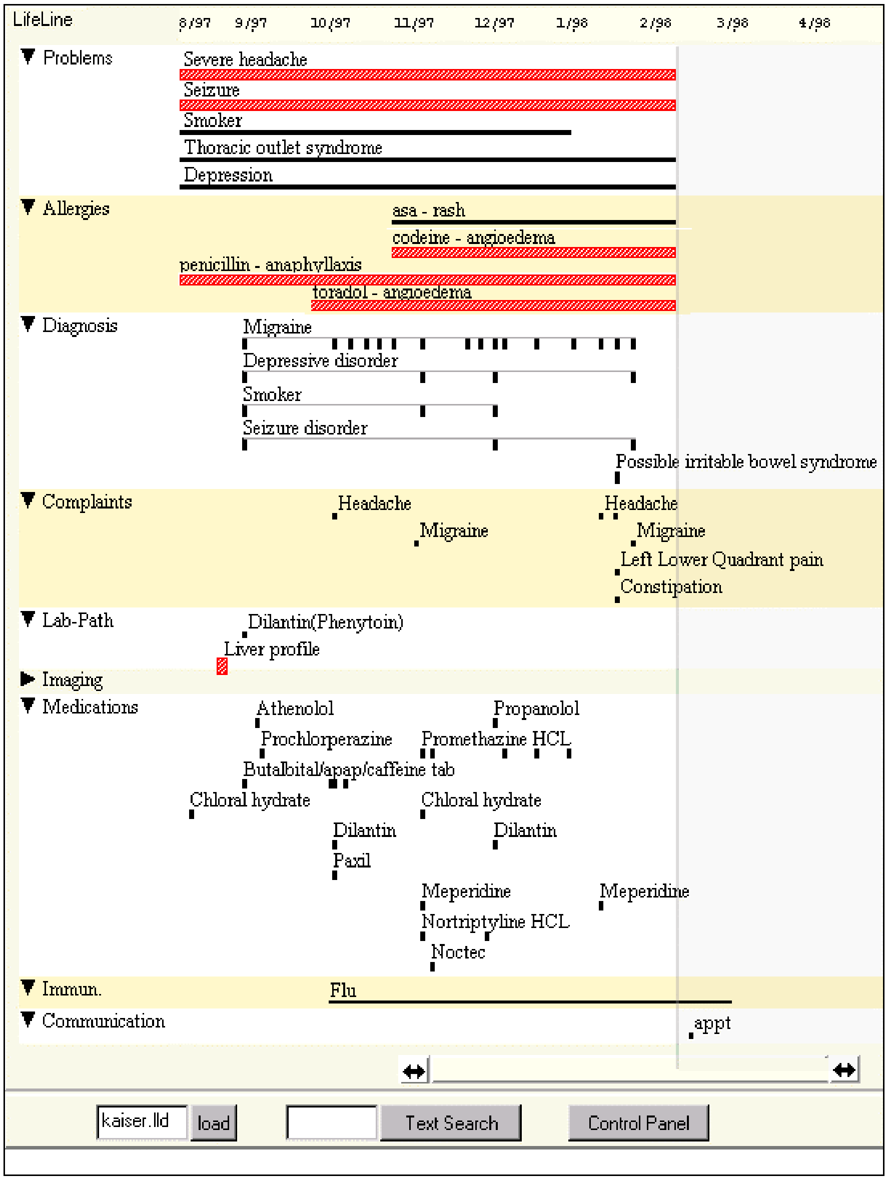

Figure 8: Example of Lifelines for the display of a patient record. Here 6 months of data grouped in facets is shown in a single screen overview. When more data is available, recursive aggregation provides a series of surrogates for the primary objects. Facets can also be closed but a silhouette preview of the facet will remains visible.

CONCLUSION

Designers of digital libraries should make the nature of their resources more apparent to users. The use of well-designed previews and overviews can help to achieve this goal. Many web sites offer rudimentary site maps, but these can be elaborated with more attribute information, including the type, size, and media of available information. Digital libraries must do even better, providing overviews that communicate the full range of data attributes common across the site so that users can see basic organizational schemes and gain a general sense of what is and what is not available in the library.

Previews should be available at a high level within a site so users get a taste of what is to come early in their visit. Exemplary samples should be selected and brought to the surface. The use of preview and overview concepts in visual information representations is clearly beneficial to the design of interactive information systems. This paper has described an information surrogate design framework in terms of what information objects are available to users (including issues related to surrogate roles, granularity, transformations, abstractness, exhaustiveness, and user attention), how information objects are related and displayed (including issues of structure such as attributes and relationships such as orderings), and how users can manipulate information objects (including issues of interactivity and tailorability). The effective design of previews and overviews results from a careful analysis of user needs and the information to be represented, the ability to determine the high-level semantics (most salient attributes) of the information, and providing directly manipulable control mechanisms that make information quickly understandable and navigable for the user. The design and testing of previews and overviews for additional types of systems and information resources will refine the framework and improve practice.

ACKNOWLEDGEMENTS

The preparation of this report was supported in part by the Library of Congress and NASA. In addition, this material is based in part upon work supported by the National Science Foundation under Grant No. IRI9615534.

REFERENCES

Ahlberg, C. and S. Truve. 1995. Exploring terra incognita in the design space of query devices. In *Proc. of Engineering for Human Computer Interaction, EHCI '95*. North Holland.

Ahlberg, C. and Shneiderman, B.. 1994. Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays , In *Proc. of ACM CHI94*, 313-317.

Bederson, B. & Hollan, J. 1994. Pad++: A zooming graphical interface for exploring alternative interface physics. In *Proc. of ACM UIST '94*. 17-26.

Borko, H. & Bernier, C. 1975. *Abstracting concepts and methods*. NY: Academic Press.

Ding, W., Marchionini, G. 7 Tse, T. 1997. Previewing video data: Browsing key frames at high rates using a video slide show interface. *Proceedings of the International Symposium on Research, development, and Practice in Digital Libraries (ISDL '97*). Tsukuba, Japan. 151-158.

Elliot, E. 1993. Watch, grab, arrange, see: Thinking with motion images via streams and collages. MSVS Thesis Document. MIT Media Lab: Cambridge, MA.

Furnas, G. W. 1997. Effective view navigation. In *Proc of ACM CHI97*, 367-374. ACM : New York.

Greene, S., Tanin, E., Plaisant, C., Shneiderman, B., Olsen, L., Major, G., Johns, S., The End of Zero-Hit Queries: Query Previews for NASA's Global Change Master Directory. In press. User Interface special issue of the *International Journal on Digital Libraries.*

Heilprin, L. 1985. Paramorphism versus homomorphism in information science. In L. Heilprin, (Ed.), *Toward foundations of information science*. White Plains, NY: Knowledge Industry Pub. pp 115-136.

Marchionini, G. 1995. *Information Seeking in Electronic Environments*. Cambridge University Press.

Marchionini, G., Plaisant, C., & Komlodi, A. In press. Interfaces and tools for the Library of Congress National Digital Library Program. *Information Processing & Management*.

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

North, C., Shneiderman, B. and Plaisant, C. 1996. User Controlled Overviews of an Image Library: A Case Study of the Visible Human. In *Proc. ACM Digital Libraries '96*, ACM: New York, 74-82.

North, C. and Shneiderman, B. 1997. A Taxonomy of Multiple Window Coordination. University of Maryland Technical Report CS-TR-3854 , UMIACS-TR-9783. http://www.cs.umd.edu/hcil/Research/tech-report-list.html.

Plaisant, C., Carr, D., Shneiderman, B. 1994. Image browsers: Taxonomy, guidelines, and informal specifications, *IEEE Software* 12:2, March 1995, 21-32.

Plaisant, C., Milash, B., Rose, A., Widoff, S., Shneiderman, B 1995. Life Lines: Visualizing personal histories, *In Proc. ACM CHI96*. ACM: New York, 221-227, 518. (see also ACM *CHI'96 Video Program)*

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. (see also ACM *CHI'98 video program*.)

Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., and Shneiderman, B. (1998), LifeLines: Using visualization to enhance navigation and analysis of patient records, to appear *in Proc. of American Medical Informatics Association Conference,* Nov. 1998, AMIA, Bethesda, MD.

Plaisant, C., Shneiderman, B., Doan, K. In Press. Interface and Data Architecture for Query Preview in Networked Information Systems, to appear in *ACM TOIS* as a Practice and Experience Paper

Plaisant, C., Shneiderman, B., Mushlin, R., An information architecture to support the visualization of personal histories, to appear in *Information Processing and Management*.

Rao, Ramana, and S.K. Card. 1994. The Table Lens: Merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. In *Proc. of ACM CHI94*, 318-322.

Shneiderman, B. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. of 1996 IEEE Symposium on Visual Languages*, 336-343.

Slaughter, L., Shneiderman, B., & Marchionini, G. 1997. Comprehension and object recognition capabilities for presentations of simultaneous video key frame surrogates. *Research and Advanced Technology for Digital Libraries, Proceedings of the First European Conference,* ECDL '97. Berlin: Springer, 39-54.

Tanin, E., Beigel, R., and Shneiderman, B. 1996. Incremental Data Structures and Algorithms for Dynamic Query Interfaces. Workshop on New Paradigms in Information Visualization and Manipulation, Fifth ACM International Conference on Information and Knowledge Management (CIKM '96) (Rockville, MD, Nov. 16, 1996) 12-15. Also in *SIGMOD Record*, Vol. 25, No. 4 (21-24).

Tanin, E, Lotem, A., Haddadin, I., Shneiderman, B., Plaisant, C., Slaughter, L. 1998. Facilitating Network Data Exploration with Query Previews: A Study of User Performance and Preference University of Maryland Technical Report CS-TR-3879, UMIACS-TR-98-14. http://www.cs.umd.edu/hcil/Research/tech-report-list.html.

Teodosio, L. & Bender,. W. 1993. Salient video stills: Content and context preserved. In *Proc. of ACM Multimedia 93*, ACM: New York, 39-46.

Tversky, A. & Kahneman, D. 1974. Judgement under uncertainty: Heuristics and biases. *Science, 185*, 1124-1131.

Tweedie, L. 1997. Characterizing interactive externalizations. In *Proc. of CHI97*, 375-382. ACM : New York.

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