

Users, User Interfaces, and Objects: Envision, a Digital Library

Edward A. Fox, Deborah Hix, Lucy T. Nowell, Dennis J. Brueni, William C. Wake, and Lenwood S. Heath
Department of Computer Science, Virginia Tech, Blacksburg, VA 24061-0106

Durgesh Rao

National Center for Software Technology, Juhu, Bombay 400049, India

Project Envision aims to build a “user-centered database from the computer science literature,” initially using the publications of the Association for Computing Machinery (ACM). Accordingly, we have interviewed potential users, as well as experts in library, information, and computer science—to understand their needs, to become aware of their perception of existing information systems, and to collect their recommendations. Design and formative usability evaluation of our interface have been based on those interviews, leading to innovative query formulation and search results screens that work well according to our usability testing. Our development of the Envision database, system software, and protocol for client-server communication builds upon work to identify and represent “objects” that will facilitate reuse and high-level communication of information from author to reader (user). All these efforts are leading not only to a usable prototype digital library but also to a set of nine principles for digital libraries, which we have tried to follow, covering issues of representation, architecture, and interfacing.

Introduction

Computer and information scientists should be among the *first* to experiment with digital libraries. In the spirit of this recommendation, the Association for Computing Machinery (ACM), as well as other associations and publishers, are becoming involved in Project Envision, a research effort supported by the National Science Foundation to build “a user-centered database from the computer science literature” (Brueni et al., 1993). Starting with users of Project Envision at Virginia Tech and spreading to Norfolk State University and other groups and individuals across the Internet, testing will proceed regarding the applicability of digital library methods to Envision’s scientific domain of computer science literature.

A goal of Project Envision is to solve some of the important research problems relating to digital libraries, especially those relating to information storage and retrieval,

human-computer interaction, and electronic publishing (Fox & Lunin, 1993). Accordingly, we have identified and tried to apply a set of principles that we believe should be the basis for future national, and later international, digital libraries. The next section explains these principles.

From the proposal stages through its current prototypes, Envision is being created as a user-centered system, as specified later in Principle 8. Therefore, users are closely involved in the development of Envision, through a structured interviewing process that guided decisions about system functionality as well as through formative usability evaluation. In the third section below, “Interviews with Users,” we discuss some of the more interesting aspects of our task analysis (Principle 7), based on user interviews. The fourth section describes the innovative Envision user interface design that evolved from this task analysis, and the results of usability evaluation of our user interface design for both the Envision query screen and search results screen.

In the fifth section, “Objects and Document Type Definition Development,” we consider how working with “objects” (see Principles 2 and 9) can help improve the overall scientific communication process, and encourage reuse of the fruits of scholarship. This has real implications regarding representation (Principles 1–3), system architecture (Principles 4–6), archiving, and use of digital library information. Finally, we conclude by highlighting some important challenges, and summarize our plans for future work.

Principles for Digital Libraries

In reviewing early work on electronic libraries, we noted the influence of current practices in traditional libraries and publishing operations. In particular, if we consider the spectrum of representations illustrated in Figure 1, we see that common practice (that is, using paper-like page images as in Elsevier’s TULIP project) may be the least useful approach for the next generation of digital libraries. Page images have all of the limitations of regular paper (problems with resizing, arbitrariness of “chunking” into pages, and so on),

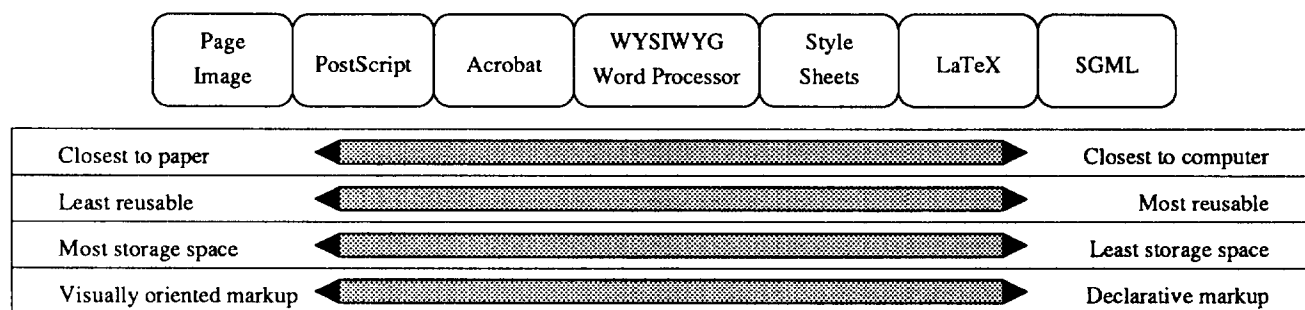


FIG. 1. Spectrum of document representations.

can only be reused by copying, require enormous amounts of storage space, consume natural resources, and are marked up in a way that virtually excludes computer exploitation of a document's organization to help in searching.

Once we go beyond our historical focus on pages, an enormous range of electronic products and services become feasible. Beginning in 1988, we worked with ACM to enlist the involvement and creativity of many individuals and groups interested in electronic publishing (Fox, 1988a). A variety of CD-ROM, hypertext, book, videotape, and online database products emerged. Then the time seemed right to take the next step, specifically, designing and building a prototype electronic archive or digital library (Fox, 1990).

Consequently, we began to reconceptualize the idea of digital libraries, as we envision their next generation. We aimed to harmonize and integrate concepts from a variety of interrelated fields: Artificial Intelligence (AI); distributed systems; electronic publishing; human-computer interaction; hypertext; hypermedia; information storage and retrieval; object-oriented approaches (analysis, databases, design, development, programming); and open systems. This focus led to the following set of nine principles for constructing electronic archives, arranged in three areas: Representation, Architecture, and User Interface. An earlier explanation of these principles was presented in Fox (1992).

Representation

Principle 1: Declarative representations of documents should be used. Linguistics and communication theory teach us to be concerned both with the content and form of documents. Document form is represented using one or more "markup" schemes, and the most usable scheme for electronic publishing is called "declarative" or "descriptive" markup (Coombs et al., 1987), which is supported by an ISO standard, the Standard Generalized Markup Language of SGML (Goldfarb, 1991). This approach lets us model documents as a collection of ordered hierarchies of content objects (OHCOS) (DeRose et al., 1990). Thus, the guidelines developed as part of the Text Encoding Initiative suggest markup conventions for

old manuscripts, poetry, dictionaries, and other literary works. Often there are multiple OHCOS, such as one for chapter and verse, and another for page and paragraph. Note that in Biblical scholarship, for example, the former, rather than the latter page-oriented, approach is preferred. Further, the many important links (see also Principle 3) inside or among documents can be flexibly captured for increased portability using the declarative ISO standard HyTime, which is a hypermedia standard based on SGML (Newcomb et al., 1991). In summary, declarative representations of documents are feasible, and standards now exist which facilitate easy interchange.

Principle 2: Document components should be represented using natural forms, namely "objects" that can be manipulated by users familiar with those objects. When we think of documents in their most general form, specifically as multimedia "bundles" of information, it becomes clear that object-oriented representations are essential. String matching systems like PAT view documents as substrings, and basic retrieval (e.g., simple Boolean or vector or probabilistic) systems concentrate on vectors of features, so it is infeasible to ask context-dependent questions or to inquire about structure as might arise in a question about inclusion relationships. As we move to multimedia documents, which are becoming more common as multimedia technologies are refined and multimedia systems become more available (Fox, 1991), the weaknesses of these models become even more evident. In particular, multiple media must be synchronized or coordinated as well as interrelated. An ISO committee, the Multimedia Hypermedia information coding Experts Group (MHEG), deals with input, output, and interaction objects and their relationships in real-time multimedia systems. Describing and processing these documents becomes so complicated that object-oriented programming, where savings arise through inheritance, is essential. User interaction is also complicated, unless various document parts each can be manipulated as a separate object: Video or audio is played or stopped; animations are run; spreadsheets are executed with new data; simulations are tried with different parameters; algorithms are executed or animated; and three-dimensional images are rotated to provide different perspectives. Mathematical

objects present unique challenges and possibilities, so that derivation or proof objects can be analyzed step by step, or formula objects can be visualized in various ways (Wolfram, 1991). Clearly, representations using objects that are convenient for users will allow authors to communicate more directly with readers, rather than going through the awkward, low-level medium of paper.

Principle 3: Links should be recorded, preserved, organized, and generalized. As we integrate documents into very large collections covering an entire scientific domain or professional area, links among those documents become increasingly important to help with search and browsing. Groupings of those links into paths, threads, tours, and webs are essential for organizing, personalizing, sharing, and preserving the structural, interpretational, and evolutionary connections that develop. We are beginning to see the emergence of wide area hypertext systems (Yankelovich, 1990) like the WorldWideWeb (WWW), that carry this concept forward into a distributed environment. Clearly, we must coordinate hypertext and hypermedia linking with the various approaches to search and retrieval (Fox et al., 1991b). One approach is the idea of information graphs (including hypergraphs), where objects of all types are interrelated by links or arcs that capture not only citation (reference) but also inheritance, inclusion, association, synchronization, sequencing, and other relationships. By specializing object-oriented databases to this task, we are building a foundation for next-generation integrated retrieval systems (Chen, 1992). Our work with the Large object-oriented External Network Database (LEND) system and methods for querying information graphs (Betrabet et al., 1993) is along these lines, as are other efforts to build systems for managing information graphs (Glyssens et al., 1990; Paredaens et al., 1992). Clearly, adaptations of hypertext (link) and semantic network (AI) concepts are essential for digital libraries.

Architecture

Principle 4: There should be a separation between the digital library and user interfaces to it. To serve millions of users, with their diversity of backgrounds, talents, and needs (see "Interviews with Users"), a variety of user interfaces will be needed for digital libraries. With hardware limitations and variations, there are a host of other reasons for building user interfaces that are particularly suited to common environments. Thus, in Project Envision, we have development efforts underway for Macintosh (specifically, both 13-inch and megapixel displays), X/Motif, and NeXT-step user interfaces. Earlier reports (Nowell & Hix, 1992, 1993a, b) and the discussion in the section on interface design below explicate these issues. With all this necessary tailoring of interfaces, it is clearly much easier if the system architecture is such that the digital library itself can be decoupled and developed separately. Common parlance

refers to a *client* system running the user interface, a *server* system managing access to the digital library itself, and a well-defined *protocol* organizing the communications required between the two. In the case of digital libraries it makes sense to begin with the Z39.50 protocol that was originally devised for communication between library catalog and bibliographic database systems. That is the approach taken in the popular Wide Area Information Server (WAIS) system (Kahle et al., 1993). We believe that further generalization is needed, so that information objects and their links can also be communicated, and we have been developing an Envision protocol to test that idea. In our case, then, we have Envision client software to manage the user interfaces, an Envision protocol, and the main (distributed) Envision system.

Principle 5: Searching should make use of advanced retrieval methods. At the heart of digital library systems like Envision, there must be support for searching, browsing, following links, presenting selected information, and other services. Regarding searching, our experimental studies, and others recently completed in connection with the 1992 Text REtrieval Conference (TREC), indicate that advanced retrieval methods can be more effective than conventional Boolean approaches. Our work with hundreds of thousands of library catalog records indicates that users prefer vector and feedback methods to standard Boolean searching (Fox, 1988b; Fox et al., 1993). These approaches can be further extended through the use of frames (Weaver et al., 1989) and other representations to get closer to "concept searching." On the efficiency side, advances in hashing (Wartik et al., 1992) can improve performance in ordered dictionaries (Fox et al., 1991a). In many indexing, linking, and other situations, guaranteed direct access to large collections, given a desired key, can be supported by rapidly finding minimal perfect hash functions (Fox et al., 1992a, b). With all these possible benefits, future digital libraries should certainly be designed to use the most advanced retrieval methods possible.

Principle 6: Open systems that include the user, and where (some of) the functions of librarians are carried out by the computer, must be developed. As digital libraries emerge, and become directly available to end-users, it is important not only to improve the user interfaces, but also to provide assistance to users like that offered by experienced librarians and search intermediaries. One approach is to develop distributed expert-based information systems, building upon studies of user-intermediary protocols (Belkin et al., 1987). Specifying the user's information need or problem, modeling the user, specifying the subject domain, and managing the overall dialog are of particular importance. Our COMposite Document Expert/extended/effective Retrieval (CODER) system was designed along these lines (Fox & France, 1987; Fox, 1987). Other efforts in this regard suggest that, while development is difficult and time-consuming, such an approach may be of value when large

numbers of users are involved. We hope that experienced intermediaries will become involved in expert system projects to pass on their guidance to millions of end-users.

User Interface

Principle 7: Task-oriented access to electronic archives must be supported. Current efforts to build prototype digital libraries are often focused on a particular subject domain, in part because of support provided by associations or publishers. Thus, the CORE project (involving the American Chemical Society and Chemical Abstracts Service, as well as Bellcore, Cornell, and OCLC) deals with the chemical literature (Lesk, 1991). Part of the hope of that project is to have access to the chemical literature be a key feature of a "chemist's workstation." Supporting the research, referencing, writing, and educational activities of staff in a university chemistry department can be viewed as providing task-oriented access to information suitable for each of those types of activities. We believe that in addition to having user interfaces that support information access as a separate activity, with its aspects of searching, browsing, previewing, and so on, "embedded information access" must be enabled. For example, a chemist preparing a class or conference presentation should be able to escape from a tool like PowerPoint™, find a description of an important reaction, grab the registry number and structure diagram for one slide, extract a table showing yield for another slide, and return directly to the expanded presentation. Similarly, a programmer accessing the Envision archive should be able to interrupt a programming effort to find a useful algorithm from *Collected Algorithms*, and add it as a subroutine, along with capturing some of its documentation and pointers to more information. We hope that efforts of this type will proceed in similar fashion to how computers in cars, microwave ovens, and compact disc players now support rather than interfere with users' tasks.

Principle 8: A user-centered development approach should be adopted. Since workstations are often devoted to individual users, we must make them serve those users. We should turn our system development efforts around to be centered on the users, rather than on the machine. Without this focus on the user, we may well produce digital libraries (and other interactive systems) that can compute perfectly and quickly, but cannot communicate effectively and efficiently with their users. As we learn more about design and development of interfaces (Hix & Hartson, 1993), a user-centered approach becomes more feasible. The next two sections explain our efforts in user-centered design of the Envision system.

Principle 9: Users should work with objects at the right level of generality. If we follow Principle 2, our digital libraries will represent information in terms of usable objects. With advanced search methods such as those called

for in Principle 5, we can search, browse, and preview those objects. Further manipulation should be supportive of user tasks, as called for in Principle 7. We consider all of these issues further in the section on objects and document type definition development.

The following sections discuss many of these principles further, focusing on users, user interfaces, and objects. For more general information on Project Envision, the reader is referred to Brueni et al. (1993).

Interviews with Users

In accordance with Principle 8, we began by focusing on potential users of a digital library of computer science literature, such as Envision. Over a four-month period we interviewed 12 professionals in the areas of computer science and information retrieval. Interviewees were chosen carefully to broadly represent the type of user we expect for Envision. During intensive interviews lasting from one to two hours, interviewees responded to questions focused on four topics:

- (1) Current information retrieval practices.
- (2) Current information dissemination practices.
- (3) Desired information retrieval and manipulation capabilities.
- (4) Demographic data.

When seeking publications relevant to a particular topic, most of our interviewees have used electronic information systems of some kind. These include computerized library catalogs, CD-ROM systems, and online search services. However, our interviewees found existing systems difficult to use for a variety of reasons. Inadequate access to any electronic information system is one major problem. Indeed, the feature most requested by interviewees for a new information retrieval system is access from the workstation in their own offices.

Interviewees also complained about the difficulty of structuring queries, the number of diverse user interfaces, inadequacy of feedback about unsuccessful searches, and the amount of knowledge required before systems are really usable. Our interviewees generally disliked any requirement or need to consult a human intermediary, or search system expert, to access the literature.

Most interviewees specifically requested or implied the need for full text retrieval. Other features commonly requested include:

- Access to multiple forms of information (abstract, resume, brief description, full text, bibliographic entry) about each document retrieved;
- Print capability;
- User annotation facilities; and
- Ability to establish and work within a personal subset of the database.

A *usable* interface was mentioned often as a needed feature, and complaints about the user interfaces of existing elec-

tronic information retrieval systems were frequently cited reasons for not using those systems.

Our interviewees want the ability to explore patterns in the literature. One spoke at length about the “community of discourse,” or invisible college of people carrying on conversations in print, all reading what the others have written. Others spoke of citation indexes, reference tools that reveal patterns of citation within the literature, so that works evolving from major articles may be identified. Ability to locate seminal documents, those which have been widely cited, is needed. Interconnections in the literature are of widespread interest. People want to use hypermedia linking to navigate among documents with common patterns of citation and to follow chains of reference among documents. In essence, they want to be able to follow on-going “conversations” in the literature.

Browsing was another common theme. Users want to be able to explore the literature along dimensions of their choosing, to home in on particular areas of interest and explore those in detail, then move on to broader views, or sometimes different views. For some, browsing includes the ability to examine the structure of documents, not just the citation or the abstract. Users want to identify a section of interest in a document and zoom in on it for closer examination and more details. Access to tables of contents provides part of this capability, but users want to move seamlessly between the table of contents and the body of a document. They want to see structure at a finer granularity than a table of contents allows. Capability to search document structures is wanted, so that chapters, illustrations, graphs, or sections of code might be located, not just whole documents by title, subject, or author.

For some, browsing is a luxury rarely permitted by pressures of time. These users want the ability to locate a few critical items of interest and be protected from the rest. They are especially interested in powerful filters to eliminate “junk” and allow them to easily locate only the most highly relevant materials. Offered the possibility of a system regularly scanning the literature for them and notifying them of new publications of probable interest, they were fearful of being overwhelmed. Information overload was cited as a reason for avoiding Internet discussion groups and bulletin boards.

Interviewees shared reliance on journals and conference attendance as major sources of information, with additional attention to conference proceedings. Talk with colleagues was ranked as equal in importance to journals as a source. Colleagues are especially helpful in providing pointers into the literature, that is, specific references to works likely to be helpful in solving a particular problem or to be of particular interest. One interviewee indicated that colleagues serve as valuable filters; they point to the few best works in an area without providing an exhaustive list of less valuable materials. A few interviewees make use of network bulletin board services, but most do not.

Interviewees indicate that they rarely use videos, because of the inability to browse or skim video, which is seen in-

stead as an “all-or-nothing” experience. Users are frustrated by the difficulty of locating particular segments of video that are of special interest. They would like to see a “video table of contents” and to be able to create hyperlinks to and from specific video frames.

We asked our interviewees about objects of interest in the computer science literature. They spoke of the obvious entities: books, journals, articles, videos, bibliographies, even figures and tables. People are also objects of interest, as authors, as researchers, as colleagues. Research projects, funding sponsors, conferences and workshops, and various types of institutions are also objects of interest. Additionally, programs, data structures, algorithms, animations, programming languages, hardware devices, interesting problems, and concepts are entities the users wish to manipulate. Users want access to source code, ideally in a choice of languages. They want to be able to embed the code in their own programs for testing and use with their own data, without rekeying the code. They would like access to analytical data about algorithms, to explanations by experts, and to animations that increase comprehension.

Design and Evaluation of the Envision User Interface

Responding to interviewees’ concern that an information retrieval system must be accessible from their offices, our design is based on the premise that the Envision user interface will run as a client process on a user’s desktop computer, communicating with the Envision retrieval system via network. Our user interface designs provide flexible use of varying configurations of monitors, both in size and number of displays. The lowest configuration supported uses a single 13-inch gray-scale display. With larger or more monitors, tiling of windows becomes feasible, and it is easier to work with full-text or page-image retrieval.

Our interface specification calls for separate windows or groups of windows for each of the major phrases or types of interaction with the Envision system. These include:

- Query Window (with four query fields and a query history);
- Search Results Windows (Graphic View, Item Summary, Item Preview); and
- Browsers.

The next two subsections deal with the Query and Search Results Windows (see Fig. 2), respectively. Work on the Browsers will be reported in a later publication.

Envision Query Window Design

The Envision Query Window design gives users the benefits of natural language query formulation (i.e., no complex syntax or use of logical operators is required, nor is knowledge of an artificial indexing language), while also providing the means to restrict searches. The Query Window has two categories of use:

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.