Digital Libraries for Creative Communities

Ian H. Witten, Matt Jones, David Bainbridge, Polly Cantlon, Sally Jo Cunningham

Department of Computer Science, University of Waikato, New Zealand

{ihw, mattj, davidb, pollyc, sallyjo}@cs.waikato.ac.nz

Abstract

Digital library technologies have a great deal to offer to creative, design[1] communities. They can enable large collections of text, images, music, video and other information object to be organized and accessed in interesting and diverse ways. Ordinary people – people not traditionally viewed as "creators" or "designers" – can now conceive, assemble, build, and disseminate new information collections. This paper explores the development rationale behind the Greenstone digital library technology. We also examine three examples of creative new techniques for accessing and presenting information in digital libraries and stress the importance of tailoring information access to support the requirements of the users and application area.

1   Introduction

Digital libraries are large, organized collections of information objects. Whereas standard library automation systems provide a digital version of the catalog—a gateway into the treasure-house of information stored in the library—digital libraries incorporate the treasure itself, namely the information objects that constitute the library's collection. Whereas standard libraries are, of necessity, important and substantial institutions, with large buildings and significant funding requirements, even large digital libraries can be lightweight. Whereas standard libraries, whose mandate includes preservation as well as access, are "conservative" by definition, with institutional infrastructure to match, digital libraries emphasize access and frequently ignore long-term preservation altogether. (Digital preservation is a thorny, and very important, subject that lies beyond this article's scope.)

As befits their role in facilitating the open distribution of knowledge, library traditions are liberal, founded on the belief that libraries should serve democracy. To help fulfill their mission as resource centers for citizens, public libraries maintain collections of records, policy statements, government documents, and so on. A recent promotional video from the American Librarian's Association declares that "the library is democracy's place of worship" (ALA, 2002).

Because of their lightweight nature and potentially far greater accessibility, digital libraries will have socially creative effects that reach even farther. Once developed, there is no requirement for institutional support. They can be distributed throughout most of the developed world over the Internet. In developing countries and remote corners of the developed world they can be circulated on removable media—CD-ROM, DVD, or gigantic disk units the size of videocassettes—and updated over radio (or in Internet cafés). Copyright is a tricky issue, but manageable. There is plenty of non-copyright material, or material whose owners are prepared to donate copyright for socially useful purposes, and trends towards more open access to academic and humanitarian information are visible. In the longer term, publishers, to remain viable, will have to work out alternative revenue models for the information they own. No wonder international organizations such as the United Nations, along with many smaller non-government organizations (NGOs), are keenly interested in digital library technology.

---

[1] For the purposes of this paper we use the term *design* as a discrete form of practice in computer science, concerned with aspects of technological functioning – in our case for the collection, structuring & accessing of information – rather than the more artistically oriented practice of design such as visual communication or graphic design.

However the *potential* for the inclusion or intervention of graphic design in the human - computer interface of a digital library is enormous and benefits could include improved accessibility and user satisfaction. The creative use by artists of digital libraries such as Greenstone also remains unexplored. These chapters have yet to be written.

Advances in technology are radically lowering the bar for the design and production of richly-organized, coherent, focused collections of information—where 'design' includes both the specification of the information structure of a digital library, and the development of the visual appearance of the interface. Now, anyone with access to sufficient source material can use public-domain software to build large, fully-searchable, collections the size of traditional personal or institutional libraries—in minutes. Let the minutes stretch to hours and the collection can be polished, organized, branded, distributed. It can include fully-illustrated text, images, video, music. It can present attractively-designed pages with consistent use of icons. Keywords, key phrases, even acronyms and their definitions, can be extracted—automatically—and used to underpin novel means of access. Let the hours stretch to days and metadata can be manually added that permits further levels of organization. Given access to programming skills, creative new facilities that stretch the imagination can be rapidly integrated into the system.

All this, one might say, can be done with ordinary Web sites: there is no need for digital library technology. However, bitter experience has shown that all but the most rudimentary sites do require significant institutional support—for organization and maintenance. The Web is littered with incomplete, unfinished, unmaintained, out-dated, inconsistently-organized, useless information collections. Just as traditional library cataloging procedures integrate new works into existing collections with minimal overhead so that they immediately become first-class members of the collection, so digital libraries allow new documents to be added completely automatically. In the case of traditional libraries this is done through the small but non-negligible overhead of generating a new catalog entry. With ordinary Web sites it requires inserting links manually into index pages and the like, and may involve editing links not only into the new document but also into existing ones that ought to reference it—it's like rewriting the book, and maybe revising all other books in the library too! Digital libraries bring access structures instantly and effortlessly up to date whenever new documents are added.

This paper argues that digital libraries enable people to exercise their creative powers to conceive, assemble, build, and disseminate new information collections. In the next section we introduce digital library technology, and specifically the Greenstone system developed at Waikato, by showing some examples. Following that, we look at the design rationale behind Greenstone. Finally, we examine three examples of creative new techniques for accessing and presenting information in digital libraries.

2   What are digital libraries?

A digital library is an organized collection of information,

> *a focused collection of digital objects, including text, video, and audio, along with methods for access and retrieval, and for selection, organization, and maintenance of the collection*
>
> (Witten and Bainbridge, 2002)

The "focused" nature of the collection is what distinguishes it from the World-Wide Web. Note the broad interpretation of "digital objects"—not just text. Beyond audio and video we also want to include such things as 3D objects, simulations, dynamic visualizations, virtual-reality worlds. The definition deliberately accords equal weight to user (access and retrieval) and librarian (selection, organization and maintenance). The latter functions are often overlooked by digital library proponents, who generally work from a technology perspective rather than from the viewpoint of library or information science, but it is precisely these aspects that allow digital libraries to be used to address social issues. Now anybody can build libraries, and give the people access.

The Greenstone software is a freely distributed, open-source product whose aim is to develop novel digital library technologies and make them available for others to use. Greenstone digital libraries are arranged in *collections*. A collection comprises several (typically several thousand, or several million) documents, and a library may include several collections, each organized differently. Collections built with Greenstone offer effective, attractive searching and browsing facilities based on metadata and the full text of electronic documents.

2.1.1   Browsing and searching for information

The most common use of Greenstone is to look for information in a collection. Most collections support both searching and browsing, although they differ depending on the collection design and the metadata available. Typically you can search for particular words that appear in the text, or within a section of a document, or within a title or section heading. A variety of interfaces exist for browsing collections by *title*, *subject*, *date*, or any other
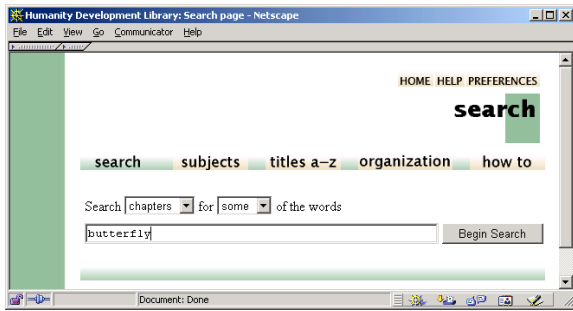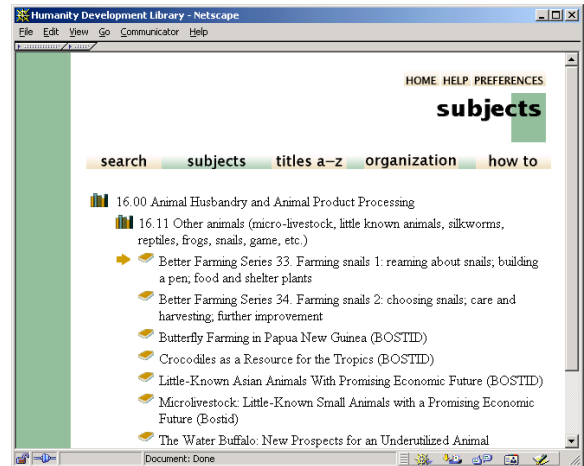
Figure 1. Searching for *butterfly*.

Figure 2. Browsing by *subject*.

metadata chosen by the collection maintainer.

In Figure 1 documents in the Humanity Development Library are being searched for chapters containing the word *butterfly*. In Figure 2 the same collection is being browsed by subject: by clicking on the bookshelf icons the user has discovered an item under *Section 16, Animal Husbandry*. Pursuing an interest in butterfly farming, the user selects a book by clicking on its book icon (Figure 3).

Documents are presented as Web pages generated by Greenstone from the source material. In Figure 3 the book's cover is displayed as a graphic on the left, and an automatically constructed table of contents appears at the start of the document. The current focus, *Introduction and Summary*, is shown in bold in the table of contents with its text starting further down the page. Greenstone collections present documents as automatically-generated Web pages. This allows documents in different source formats to be presented in a consistent manner, and lets users view the entire collection with a standard Web browser—no special viewing applications are required. However, the collection maintainer may choose to present the original source document instead of, or as well as, the HTML version.

All the icons in the screenshots of Figures 1-3 are clickable. Those at the top of the page return to the library home page, provide help text, and allow you to set user interface and searching preferences. The navigation bar beneath gives access to the searching and browsing facilities, which differ from one collection to another.

2.1.2 Internationalization

The international Unicode character set is used throughout, so that documents in any language and character encoding can be imported. (In fact, Greenstone can automatically detect the language and encoding of most documents.) Collections of documents in Arabic, Chinese, Cyrillic, English, French, Spanish, German, Hindi, and Maori are publicly available. The New Zealand Digital Library Web site (nzdl.org) hosts many of these, and the Greenstone Web site (greenstone.org) links to sites that contain further examples.

It makes little sense to have a collection whose content is in Chinese or Russian, but whose supporting text—instructions, navigation buttons, labels, images, help text, and so on—are in English. Consequently, the entire Greenstone interface has been translated into a range of languages, and the interface language can be changed by the user as they browse from the *Preferences* page. This is possible because the language-dependent images and Web pages are generated from a set of "macro files" that have been translated by Greenstone users in other parts of the world and contributed back to the project. (The same mechanism provides text-only versions of the interface to accommodate visually impaired users.) Figure 4 shows an example: a Russian collection.
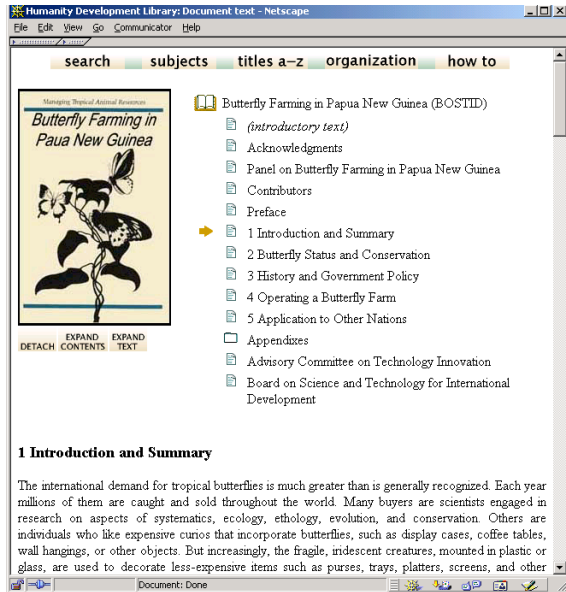
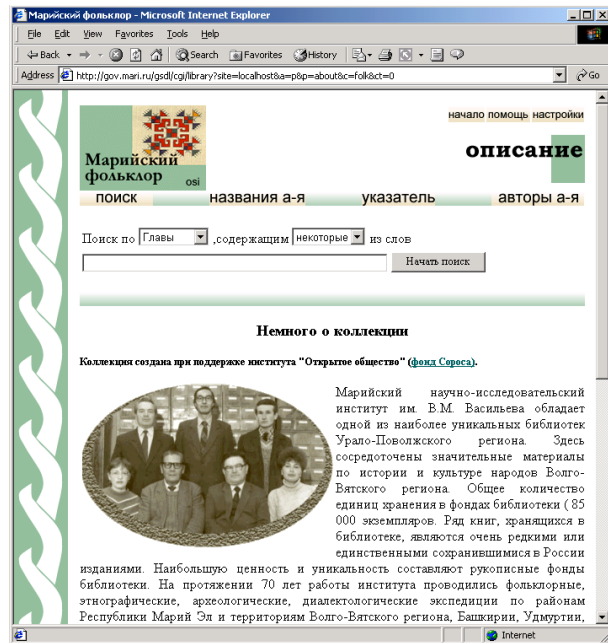Figure 3.   A document: *Butterfly farming in Papua New Guinea*.

Figure 4.   A Russian digital library collection

3   Greenstone's Rationale

Greenstone has been developed over several years in a collaborative, distributed way. Although key developers have been with the project since its inception, many contributors have worked for 1-3 years as part of the team, before moving to other projects. Many members of the team have been based at Waikato, but the open-source nature of the work (described in Section 3.1) has led to helpful efforts from around the world.

To help guide this rich, complex design environment, the project has a number of guiding design principles that include:

anyone can contribute to the design and development of the system;

many infrastructures should be supported;

anyone can build a collection; and,

anyone can benefit from a digital library

3.1   Anyone can contribute to the design and development of the system

Greenstone is open-source software, issued under the terms of the GNU General Public License[2]. The GNU organization identify five "freedoms" software users should be entitled to:

---

[2] http://www.gnu.org/licenses/gpl.html

"The freedom to run the program, for any purpose (freedom 0)

The freedom to study how the program works, and adapt it to your needs (freedom 1). Access to the source code is a precondition for this.

The freedom to redistribute copies so you can help your neighbor (freedom 2).

The freedom to improve the program, and release your improvements to the public, so that the whole community benefits (freedom 3). Access to the source code is a precondition for this." (From www.gnu.org)

By subscribing to this philosophy, we aim to encourage anyone to adapt existing functions, add new features and apply our software in ways we could never have envisaged. While programming skills are necessary to extend the functionality of Greenstone proper, we are exploring development options that will allow those with limited or no programming skills to customize an existing digital library's interface and structure. The open-source approach, then, provides a highly creative, collaborative and global context for further technological developments.

3.2 Many infrastructures should be supported

Not all users have access to the highest specification, latest computer technology. In the developing world, for instance, while computers *per se* are not so rare as one might think, access to the latest models is uncommon. The extremely rapid obsolescence of personal computers combined with the developed world's voracious appetite for the latest and greatest, makes low-end machines essentially free: many find their way abroad, donated to various aid organizations.

In developing Greenstone, we have been keen to find ways of providing its facilities to users of older computing infrastructures, such as Windows 3.1. In addition, recognizing there is the serious lack of access to high-speed networks in many parts of the world, the system can be used to create CD-Rom based collections. For example, consider the Humanity Development Library, a compendium of some 1,200 authoritative books and periodicals, produced by many disparate organizations—UN agencies and other international organizations—on various areas of human development, from agricultural practice to economic policies, from water and sanitation to society and culture, from education to manufacturing, from disaster mitigation to micro-enterprises. It contains 160,000 pages and 30,000 images, which if printed would weigh 340 kg, cost $20,000, and occupy a small library book stack. Instead, it takes the form of a digital library and is distributed on a CD-ROM throughout the developing world at essentially no cost. (It's on the Web at nzdl.org.)

Recently we have been investigating the use of digital library collections on small screen devices such as mobile phones and handheld PCs (Buchanan *et al* 2002; Marsden *et al* 2002). These sorts of device are attractive both because of their likely universal pervasiveness and the fast-paced growth of wireless, mobile networks in developing countries.

Although the developing world has driven such design efforts, other groups are likely to benefit – for instance, small community groups with limited access to hi-end resources are able to create useful collections.

3.3 Anyone can build a digital library

New interfaces in the Greenstone suite make it easier for anyone to collect information into a digital library and distribute it. Figure 5 shows one such tool, known as the Gatherer, being used to collate a selection of images for a digital library collection, augment these source documents with textual metadata and then build and view the collection. From here, it is a matter of a few further clicks to produce a self-installing CD-ROM version of the collection. The figures show a user developing a digital library collection of historic paintings of New Zealand.

The user initiates a new collection through the file menu (Figure 5a). The resulting popup window prompts for some preparatory information: the title of the collection, a short directory name that represents the collection on the file system, a contact e-mail address for the person(s) maintaining the collection, the types of metadata terms that will be used to describe the source documents, and a statement describing the contents of the collection.

Once the user has filled out this form the main window becomes active. A series of menu items, *Gather*, *Enrich*,
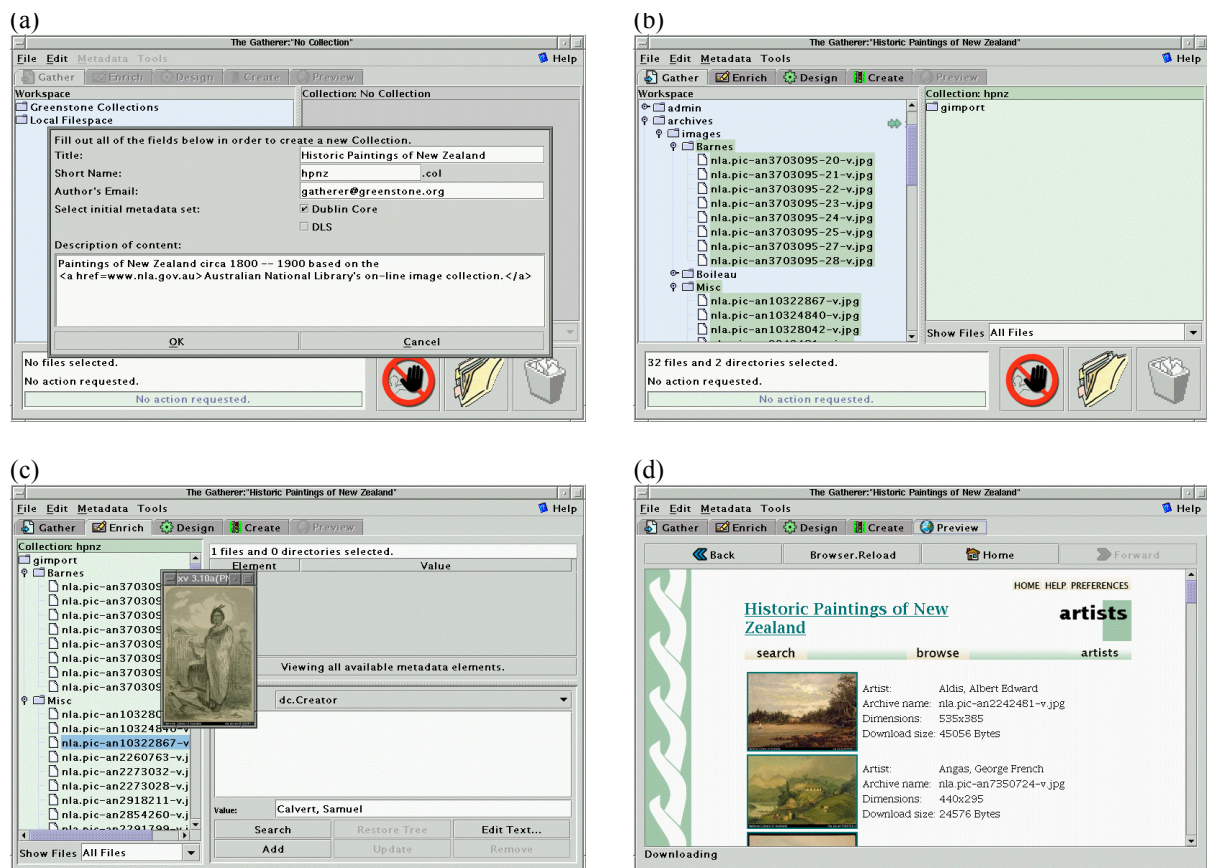
Figure 7.   Building a collection with the Gatherer.

Design, *Create*, and *Preview*, guide the user through processes involved in building the collection. Using the left-hand pane of the initial *Gather* panel, shown in Figure 5b, the user browses the file system for suitable images. The user has opened the *archives* directory, selected subdirectories *Barnes* and *Misc*, and is in the process of importing them into the digital library collection by dragging and dropping the selected files from the left pane to the right pane. This activity can be repeated as many times as is necessary. The *Boileau* directory was omitted because closer inspection by the user showed that it contains black and white photographs, not paintings.

Users typically proceed left to right through these panels. In Figure 5c they move to the *Enrich* panel and add textual metadata (the name of the artists) to the selected documents. When the collection was initialized, the user selected Dublin Core—a popular web-based metadata standard—as the type of metadata used. Here the user is entering the value "Calvert, Samuel" for the *dc.Creator* metadata term, which is the appropriate Dublin Core category for representing artists. This activity can be repeated as often as is necessary. For the paintings by Barnes, the artist's name can be applied in one operation by performing a group select.

The next two panels, which are omitted for brevity, help the user structure the collection, control its appearance, and build it. In the *Design* panel this collection was configured to produce an alphabetical list of artists (based on metadata) and customized to display source documents as thumbnail images accompanied by bibliographic information. The *Create* panel activated the building process. On completion the result is viewed in the *Preview* panel.

Figure 5d shows a page from the newly built collection, in which source documents are alphabetically listed by artist. Shown alongside each thumbnail is the artist's name, its catalog number, image dimensions, and its download size. The full-size image is shown by clicking on the thumbnail. In the Gatherer, users may pass backwards and forwards through the panels using them to augment and enhance the collection, perhaps adding in further source documents, editing metadata values, and altering the appearance of the collection.Anyone can benefit from a digital library

3.4   Anyone can benefit from digital libraries

Traditional libraries have been seen, perhaps, as somewhat "elitist" institutions: they are places that literate, advantaged people go to further their studies. Digital libraries can overthrow such dated conceptions, opening up access and creativity to new constituents. We have been exploring such possibilities and report two examples, below.

6

3.4.1 Digital libraries for oral cultures

Digital libraries have the potential to bring information to the illiterate—though we prefer to think of the target audience as people from predominantly oral cultures, and people who are literate but not in the language in question. To take just one example, consider the disaster relief scenario. Populations engaged in responding to and mitigating disasters in developing countries are usually illiterate or semi-literate, and have rarely received appropriate training. According to the UN's *Human Development Report* (UNDP, 1999), 20% of the world's population is illiterate, and the rate is twice that in Sub-Saharan Africa, the Middle East, and Southern Asia—with even higher figures for the female population.[3]

A multidisciplinary team from the Department of Computer Science at Waikato University, New Zealand and the Payson Center for International Development and Technology Transfer of Tulane University, New Orleans, together with aid workers in Costa Rica, are investigating the use of information systems by illiterate and semi-literate populations for disaster preparedness, response and mitigation in the developing world. The very notion of digital libraries for the illiterate is almost paradoxical. However, the potential payoff is enormous in terms of human development.
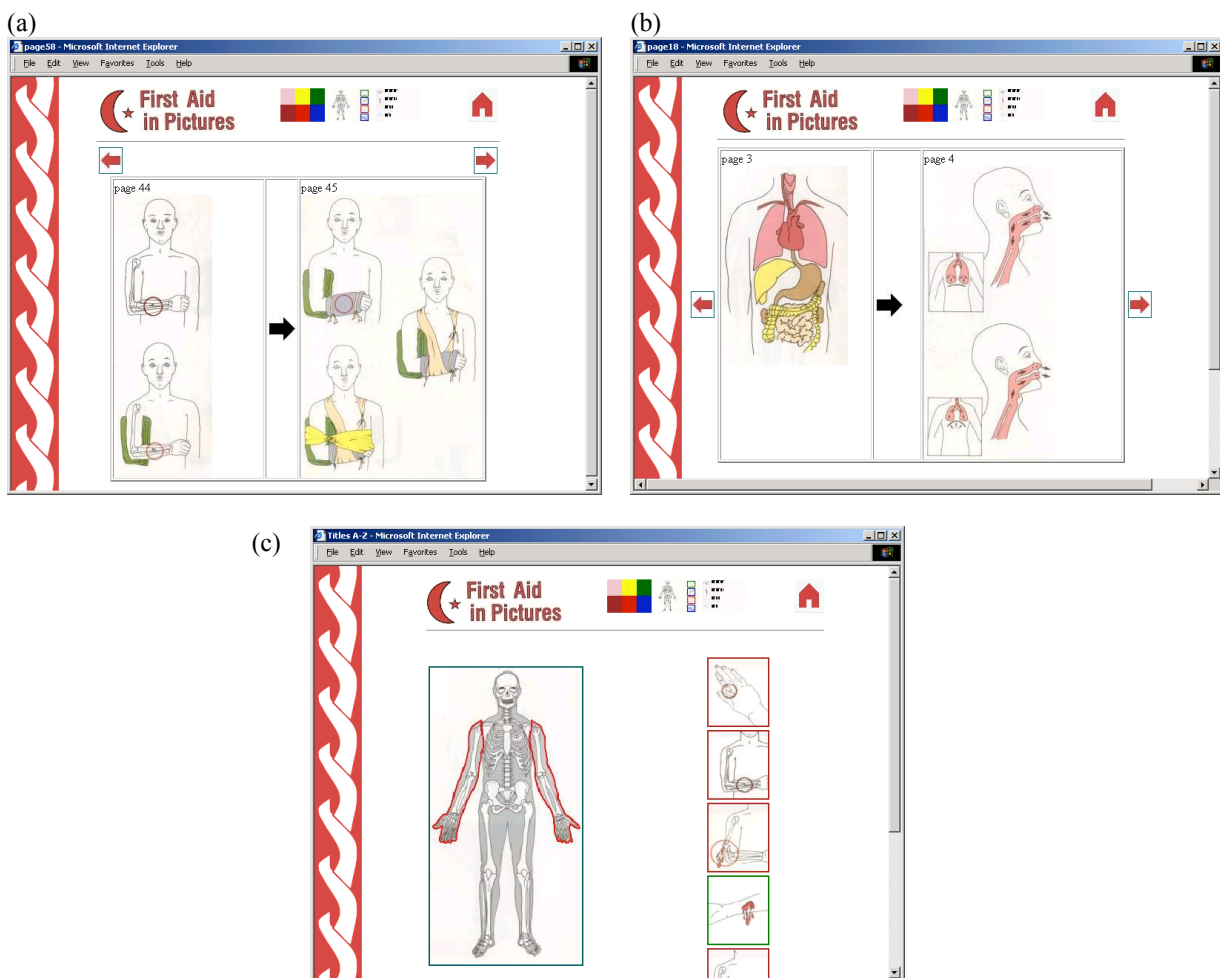

Figure 6. *First aid in pictures*: a prototype digital library collection without words

Figure 6 shows some examples from a prototype collection of first aid information. The source material is *First aid in pictures* by DeLong *et al.* (1987), a 120-page manual on first aid that contains not a single written word. Figure 6a shows a fracture of the lower arm, and how to keep the injured part from moving and hold it in place while the patient is taken for treatment by someone with appropriate medical training. Literate as you

---

[3]A 1993 survey found an illiteracy rate of 20–23% in the US (Kirsch *et al.*, 1993), though the figures are open to various interpretations—for example, an optimistic 5% (Baron, 2002).

undoubtedly are, in an emergency even you might prefer these pictures—which in the book are vividly colored, adding further clarity—to a textual account of how to deal with a medial radius fracture. As well as instructions specific to particular injuries, the collection also contains general medical information. For example, on the left in Figure 6b are five major organs of the body: lungs, heart, liver, stomach, and intestines, while on the right is the body's airway system.

Arranging convenient access to such an information collection challenges traditional notions of "searching" and "browsing." Clicking on the skeleton Figure 6c brings up a pictorial menu of thumbnails of relevant injuries; clicking on one of these takes you to appropriate target information (for example, the second one takes you to Figure 6a). This system is currently in an experimental state: much research is needed to tailor it to particular target audiences. One possible scenario is to use it for teaching, with a literate instructor who can switch on textual descriptions beneath each image. The participants, presumed illiterate or semi-literate, can take the digital library collection away after the course as a reference guide for practical emergencies.

3.4.2 Digital libraries for the disadvantaged

We have been developing a digital library for a user, "Jane" (a pseudonym), who has a profound physical disability that means she cannot communicate verbally, and cannot use conventional communication tools (Dunlop *et al.,* 2002). We used several techniques to identify expressions for the digital library and then prototyped interaction schemes to enable Jane to browse and search for things to "say".

Jane's input device is a head switch, which restricts her input events to, effectively, a single left mouse button click. Given this forced choice of input device, the only selection algorithm available is scanning: onscreen options are successively highlighted until she selects one option with a switch action.

To use browsing to select an expression, Jane first selects *Browse* mode. She is then presented with a list of the eight top-level expression metadata categories, and the scanning algorithm iterates until she either selects one of the categories, or selects the option returning her to the main screen. This process is repeated until the set of expressions in a category are displayed, and the list is scanned until she selects one of the expressions.

In order to search for an expression in the digital library, Jane must be able to enter a keyword into the digital library search interface. The list of expressions containing that term is then displayed in a conventional "hits list" screen, and the hits are successively scanned until she selects the desired expression.

In the prototype, the system was used to allow Jane to communicate face-to-face with another person. However, she already enjoys surfing the Web. If the output of the communication aid was piped to an Internet Chat session, the user could participate in conversations with others in remote locations. The relatively slow conversational speed achieved with this tool would be to some extent obviated by the delays common in Internet Relay Chat—in this case, helping Jane's conversation more closely approach social norms.

4 Creative access and presentation

Providing effective and comfortable ways to work with the huge volumes of information that are present in even modest-sized digital libraries is a continual challenge. In response, Greenstone supports creativity in the design of new interfaces to digital libraries—enabling new ways of accessing and presenting media. The digital library creator is not forced to to use a 'vanilla' searching or browsing interface, but instead can create user support structures tailored to the needs of the digital library's target user group. In this section we look at three examples for text, audio, and image material respectively: hierarchical phrase browsers formed from the full text of information collections, "query-by-humming" retrieval of musical melodies, and collage-based browsing of images.

4.1 Phrase browsing

Naturally, people often want to browse information collections based on their subject matter. This is conventionally supported by displays based on hierarchical classification metadata associated with each document. But manual classification is expensive and tedious for large document collections. If this information is unavailable, an attractive alternative is to use phrase metadata that has been extracted automatically from the
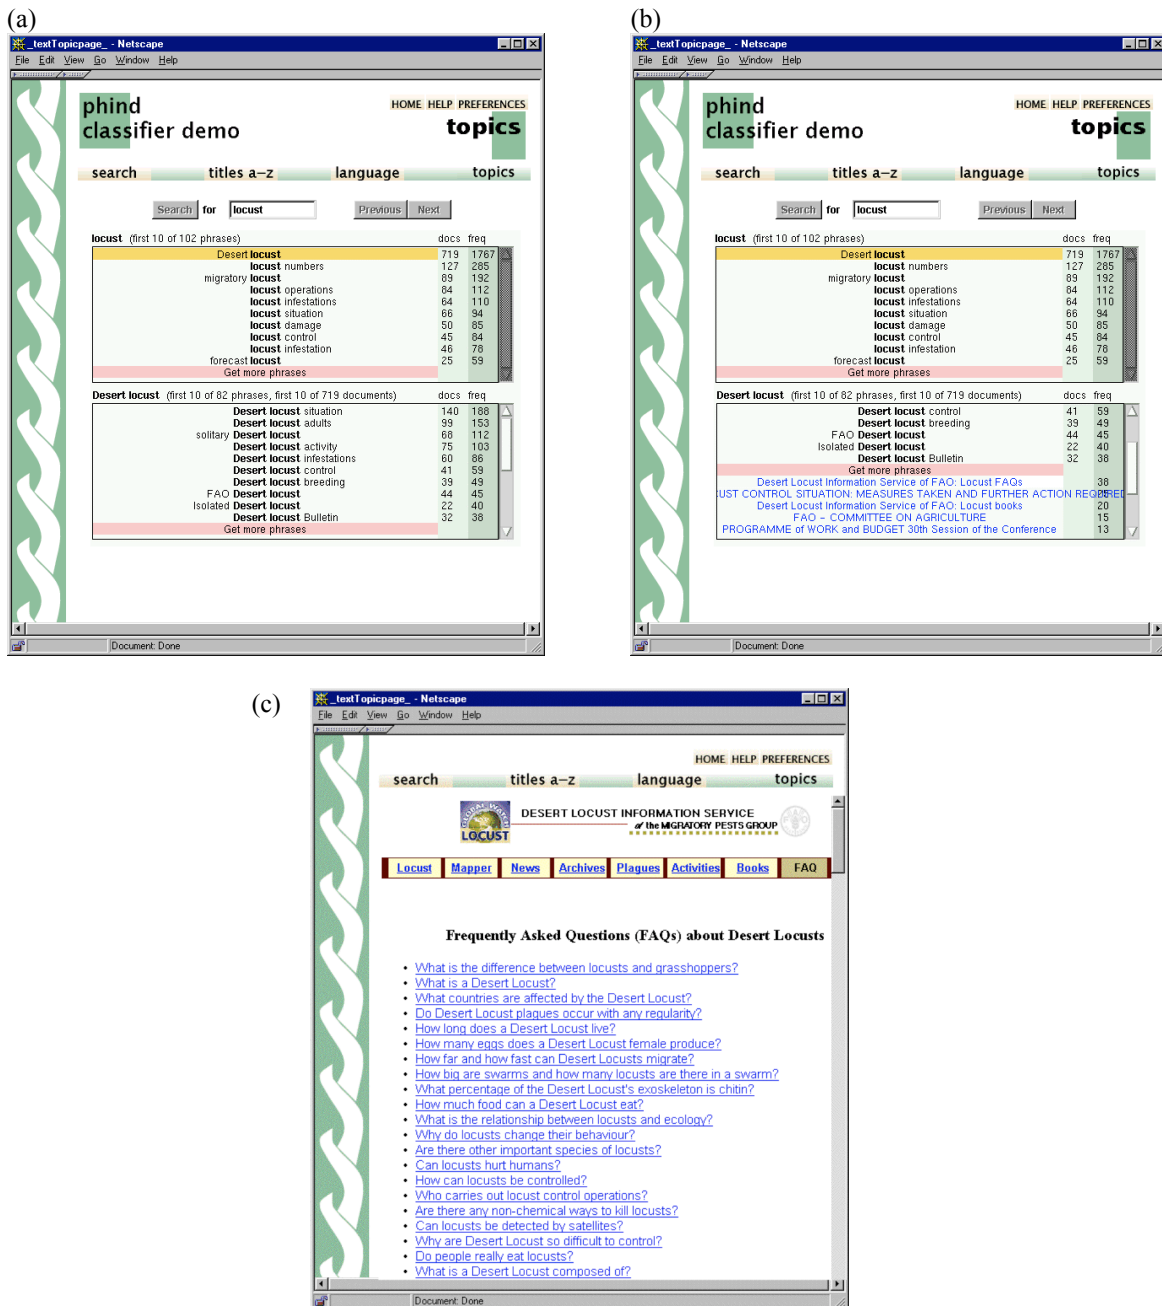
(a)

(b)

(c)

Figure 7. (a) Browsing for information about *locusts*, (b) expanding on *desert locust*, (c) document about *desert locusts*

full text of the documents themselves (Paynter *et al*., 2000).

Figure 7 shows an interactive interface to a phrase hierarchy that has been extracted automatically from the full text of a document collection. In this case the collection is simply a web site: the site of the United Nations Food and Agriculture Organization (FAO). Figure 7c shows a typical target document—a Web page, complete with fully-functional internal and external links. This interface is designed to resemble a paper-based subject index or thesaurus. In Figure 7a, the user enters an initial word in the search box at the top—in this case the word *locust*. On pressing the *Search* button the upper of the two panels appears. This shows the phrases at the top level in the hierarchy that contain the search term. The list is sorted by phrase frequency; on the right is the number of times the phrase appears in the entire document collection, and beside that is the number of documents in which the phrase appears.

9

Only the first ten phrases are shown to avoid downloading a large number of phrases—many of these phrase lists are huge. At the end of the list is an item that reads *Get more phrases* (displayed in a distinctive color); clicking this will download another ten phrases, and so on. The interface accumulates the phrases: a scroll bar appears to the right for use when more than ten phrases are displayed. The number of phrases appears above the list: in this case there are 102 top-level phrases that contain the term *locust*.

So far we have only described the upper panel in Figure 7a. The lower one appears as soon as the user clicks one of the phrases in the upper list. In this case the user has clicked *Desert locust* (which is why the first line is highlighted in the upper panel), causing phrases containing the string *Desert locust* to be displayed in the lower panel.

If you continue to descend through the phrase hierarchy, eventually the leaves will be reached. In this system, any sequence of words is a "phrase" if it appears more than once in the entire document collection. Thus a leaf corresponds to a phrase that occurs in a unique context somewhere in the collection (though the document that contains that contextually unique occurrence may include several other occurrences too). In this case, the text above the lower panel shows that the phrase *Desert locust* appears in 82 longer phrases, and also in 719 documents. These 719 documents each contain the phrase in some unique context. The first ten documents are visible when the list is scrolled down, as is shown in Figure 8b.

In effect, both panels show a phrase list followed by a document list. Either list may be empty: in fact, the document list is empty in the upper panel because every context in which the word *locust* occurs appears more than once in the collection. The document list displays the titles of the documents. In both panels of Figures 8a and b you can click *Get more phrases* to increase the number of phrases that are shown in the list. In the lower panels you can also click *Get more documents* (again it is displayed at the end of the list in a distinctive color, but to see that entry you must scroll the panel down a little more) to increase the number of documents that are shown.

Clicking on a phrase expands that phrase. The page holds only two panels, and if a phrase in the lower one is clicked the contents of that panel move up to the top to make space for the phrase's expansion. Alternatively, clicking on a document opens that document in a new window. In fact, the user in Figure 7b has clicked on *Desert Locust Information Service of FAO: Locust FAQs*, and this brings up the page shown in Figure 7c. As Figure 7b indicates, that document contains 38 occurrences of the phrase *Desert locust*.

Utilizing phrases extracted automatically from the document text, as these examples are, means that no manual processing is needed to generate these indexes. This phrase interface also has the advantage of being readily comprehensible by users. The notion of browsing documents based on phrases that they contain has great intuitive appeal, because you are in direct contact with the raw material of the documents themselves, without any intermediary operations that you may only dimly understand. From a digital creativity perspective, this type of information exploration may appeal to those working with collections of literature, where patterns of language usage are of interest.

An important feature of this interface is that it scales up to large collections. Other browsing techniques like alphabetic lists of titles and authors do not scale well: ranges of letters are uncomfortable to use for large collections. Hierarchical classification schemes scale well, but require manual classification of each document. This phrase interface is easy to browse even for colossal collections, and users are not overwhelmed with information, nor with difficult choices. It is unique in being both scalable and automatically generated.

4.2 Music retrieval

Digital libraries have the potential to revolutionize the ways in which we access and listen to music. Apple's iPod, which allows one to create a personalized collection of music and download it into a mobile device, has been called the world's first consumer-oriented digital library. More fundamentally, new techniques of melody matching provide completely new ways to access music collections that transcend what is possible in conventional music libraries (McNab *et al.*, 1996; Bainbridge *et al.*, 1999).

The user in Figure 8 is searching for a tune by humming, whistling, or singing a snatch of it (though without words)—or entering it on a music or computer keyboard. The system is capable of interpreting audio input as a sequence of musical notes and searching for that sequence in a database of melodies. The first stage in the process is to transcribe the acoustic query into symbolic musical notes. The background window in Figure 8 shows, on the left, the query in symbolic form. This has been produced by the computer from audio input
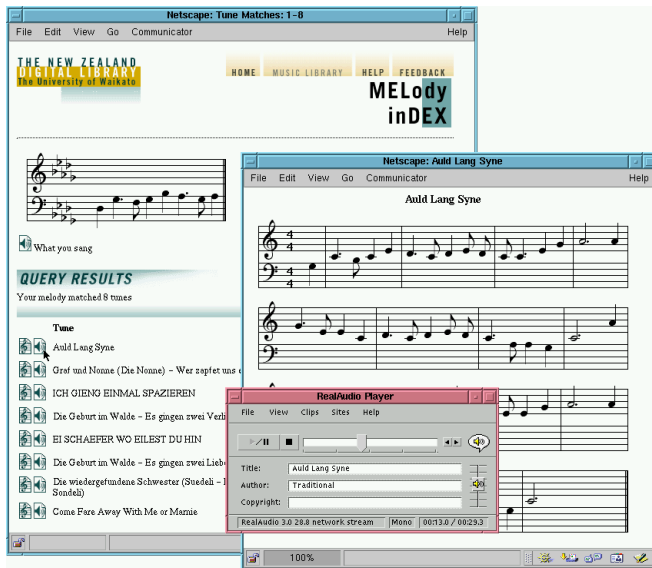
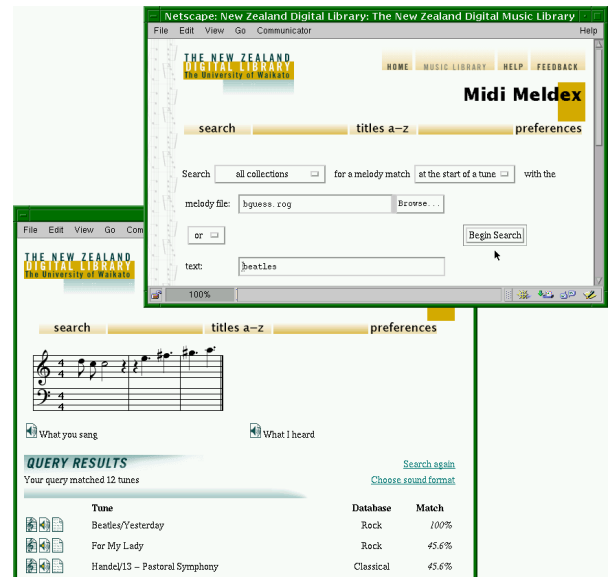Figure 8. Finding *Auld Lang Syne* in a digital music library

Figure 9. Combined music and text search

captured by a special web browser plug-in when the user sang into a microphone. The frequency profile of the input is analyzed by standard pitch tracking software, and the notes are segmented based on the amplitude profile to generate the musical representation shown. In fact, the notes are a perfect transcription of the opening bars of *Auld Lang Syne*, rendered in a nice tenor voice. Buttons underneath allow the user to replay his vocal input ("What you sang," visible in Figure 10) and play a computer synthesis of the transcription it has made ("What I heard," not visible).

The next step involves searching through a database of melodies, measuring the similarity of the input to each melody. In this case the database contains 10,000 folk tunes from around the world. The technical details of the matching operation are beyond our scope here: suffice it to say that an approximate string matching algorithm based on dynamic programming is used to determine the degree of similarity of each melody to the input. Error-tolerant matching is necessary for several reasons:

the input is noisy: pitch tracking can fail and people often sing badly;
people do not remember tunes exactly;
most melodies have different versions—in particular, rhythmic variations;
the database may contain errors.

The effect of this step is to identify melodies in the database that match the query approximately, and sort them by degree of match. The result is shown in the "Query results" list in the lower half of the background window in Figure 9. In this case the correct tune, *Auld Lang Syne*, is the best match.

Text searches can be combined with melody matching to yield a more comprehensive search capability. Figure 9 shows the query page for a collection of MIDI (Musical Instrument Digital Interface) tunes, in which a textual search for the word *beatles* is being combined with a sung melody that resembles the first new notes of the tune *Yesterday*. The music displayed in the Figure is the computer's rendition of the user's singing—note incidentally that the rhythm of the notes is disturbed because the output module, which resynthesizes the music notation into an image for display, has assumed, incorrectly in this case, that the tune starts at the beginning of a bar. This mis-interpretation does not affect the result of melody matching.

Text matching is governed by options specified on the Preferences page—here matching is case insensitive with stemming disabled. Likewise, music matching is also governed by options. Here it is restricted to the start of each tune (ignoring leading rests), takes account of the interval between notes (rather than using simply the up-down-same "contour," which is more appropriate for poor singers), and ignores note duration. From the Query results page, part of which is shown in Figure 9, items in the collection can be viewed in various forms, symbolized using icons to the left: the MIDI file reconstructed in the form of sheet music, audio playback of the MIDI file, and an HTML page showing the text that accompanies the song.

Figure 10   Collages created by the Collage machine: (a) a news collage; (b) result of a search on *archeology*

Music information retrieval involves many novel techniques. For example, it is useful to extract musical motifs—short sequences of notes that are repeated throughout a given tune. Motifs in music are analogous to key phrases in text, and new techniques of phrase extraction can be used to identify candidates. Coming up with an algorithm to identify just those motifs that are musically interesting is a challenging research problem. However,

extracted motifs are useful even if they are diluted by many uninteresting ones. For example, searching can be accelerated by producing an index of motifs in advance, like the full-text indexes used for text. Motifs could form the basis of a musical browser, serve as brief surrogates for tunes, and even underpin musicological studies. Other problems are presented by polyphonic matching—most music comes in parallel streams of information played by different instruments but strongly coordinated in time—and rhythmic matching—as any dancer will confirm, music is characterized by different rhythmic patterns.

4.3   Collaging

The word "collage" derives from the French *collé*, which means pasting or gluing, and traditional collages are static objects, painstakingly constructed. Computer technology introduces radical new possibilities. It invites a temporal element: dynamic collages. It can adjust the images' transparency to create a three-dimensional layering effect rather than a two-dimensional juxtaposition. And, most importantly, it allows collages to be interactive objects that respond to the user's actions.

Kerne (2000) has developed a "Collage machine" that presents a dynamic collage of images and snippets of text from the Web. Of course, computer-created collages lack the artist's understanding of the content of images and the emotive force of nonstandard juxtapositions. However, Kerne uses the meager context that the Web provides—the fact that some images are related because they are included in the same document, that documents are related both because they reside on the same site and through explicit linkages—to affect which images are displayed next and whereabouts on the screen. He uses the temporal and transparency dimensions to make images fade away over time. He allows users to point to images, bring them to the front, and explicitly indicate that certain ones are to be favored or deprecated. These factors are combined into a "weight" which affects the random process by which new images are chosen for display and positioned on the screen.

Figure 10 shows examples created by the Collage machine (from Kerne, 2000). Figure 10a is a news collage which pulls content from *The New York Times*, *CNN*, *USA Today*, and the *BBC*. Figure 10b is constructed from a Web search on *archeology*. These views show the screen after the Collage machine has been working for several minutes, and change dynamically at a comfortable viewing rate. Each collage is a dynamic and unique experience, being driven by random choices that are affected by the viewer's interactions.

Collages provide a radical new way of browsing around large collections of information, and are likely to be important mechanisms of information access in future digital libraries. They are particularly suited to visual imagery, but work well for text snippets too: the Collage machine selects small pieces of text and renders them for display amongst the images. Collage browsing is a relatively passive activity, more like wave watching than active surfing. Whereas conventional searching and browsing requires active engagement in which users are constantly expected to make decisions, collage browsing provides a calmer, more relaxing, and perhaps more

reflective experience which does not require the user's active participation—but responds if it occurs.

5 Conclusions

A library, by its very nature, is both developed by a community and exists to support the activities of a community of users. Digital libraries have great potential to enable communities to further influence and extend their audiences.

Novel interface technologies are opening up new ways of accessing and presenting information. Good, creative, design surely benefits society, and good design in the service of information access yields double benefit. For example, interactive melody matching is a striking example of a search paradigm that differs radically from the standard full-text search. Analogously, hierarchical phrase-based browsing of text is a radically new access mechanism for traditional information collections, and spatio-temporal collages of images and textual snippets perform a similar service for images. Perhaps these creative new interfaces herald a calmer, cooler, more reflective approach to user interfacing than the ubiquitous, frenetic, point-and-click model. Extended periods of contemplative wave-watching are part of the attraction of real surfing, interspersed with the frenzied excitement of real-time balance and urgent microdecision-making when riding the foam itself. Periods of reflective calm between frenetic spurts of web-surfing, bathed in collaged imagery, may engage different parts of our brains and provide new and more productive ways of dealing with today's information tsunami.

Opening up digital libraries for the illiterate is a radical and potentially revolutionary benefit of new interface technology. Perhaps digital libraries for people from oral cultures will find a place in our own society too, and help reduce the various "digital divides" that cleave our world (Norris, 2001)—the "social divide" between the information rich and the information poor in our own nations, the "democratic divide" between those who do and do not use the panoply of digital resources to engage, mobilize and participate in public life, as well as the "global divide" that reflects the huge disparity in access to information between people in industrialized and developing societies.

References

ALA (2002) "Rediscover America @ your library." Video produced by the American Library Association, Chicago, IL. Available from www.ala.org/@yourlibrary/rediscoveramerica.

Bainbridge, D., Nevill-Manning, C., Witten, I.H., Smith, L.A. and McNab, R.J. (1999) "Towards a digital library of popular music." *Proc ACM Digital Libraries*, Berkeley, CA, pp. 161–169.

Baron, D. (2002) "Will anyone accept the good news about illiteracy?" *The Chronicle of Higher Education*, February 1, 2002.

DeLong, M.B., Brady, J.G., Bourgeois, L.D. and Niemiec, L.J. (1987) *First aid in pictures*. Vade Mecum Press, Sterling, VA.

Dunlop, H., Cunningham, S. J. & Jones, M. (2002). "A digital library of conversational expressions: helping profoundly disabled users communicate", *Proc ACM/IEEE Joint Conference on Digital Libraries*, Portland, Oregon.

Kerne, A. (2000) "CollageMachine: an interactive agent of Web recombination." *Leonardo*, Vol. 33, No. 5, pp. 347–350.

Kirsch, I.S., Jungeblut, A., Jenkins, L. and Kolstad, A. (1993) *Adult illiteracy in America: a first look at the results of the national adult literacy survey*. National Center for Education Statistics

McNab, R.J., Smith, L.A., Witten, I.H., Henderson, C.L. and Cunningham, S.J. (1996) "Towards the digital music library: tune retrieval from acoustic input." *Proc ACM Digital Libraries*, Bethesda, MD, pp. 11–18; March.

Norris, P. (2001) *Digital divide? Civic engagement, information poverty and the Internet worldwide*. Cambridge

University Press, New York

Paynter, G.W., Witten, I.H., Cunningham, S.J. and Buchanan, G. (2000) "Scalable browsing for large collections: a case study." *Proc ACM Digital Libraries*, San Antonio, TX, pp. 215–223; June.

UNDP (1999) *Human development report 1999*. UNDP/Oxford University Press, New York (p.176).

Witten, I.H., Loots, M., Trujillo, M.F. and Bainbridge, D. (2002) "The promise of digital libraries in developing countries." *The Electronic Library*, Vol. 20, No. 1, pp. 7–13.

Witten, I.H. and Bainbridge, D. (2002) *How to build a digital library*. Morgan Kaufmann, San Francisco.