

Enhanced browsing in digital libraries: three new approaches to browsing in Greenstone

Dana McKay, Preeti Shukla, Rachael Hunt, Sally Jo Cunningham

Department of Computer Science, University of Waikato, Private Bag 3105, Hamilton, New Zealand
e-mail: {dana, ps11, rmh14, sallyjo}@cs.waikato.ac.nz

Published online: 20 October 2004 – © Springer-Verlag 2004

Abstract. Browsing is part of the information-seeking process, used when information needs are ill-defined or unspecified. Browsing and searching are often interleaved during information seeking to accommodate changing awareness of information needs. Digital libraries often do not support browsing. Described here are three browsing systems created for the Greenstone digital library software.

Keywords: Browsing – Information seeking – Searching – Self-organizing maps – Thesaurus

1 Introduction

Browsing is a vital part of the information-seeking process, allowing information seekers to meet ill-defined information needs and find new information [26, 29]. Despite the importance of browsing in information seeking, however, it is poorly supported in many information systems [27]. The Greenstone digital library software [1] created by the New Zealand Digital Library research group [2] is an example of software that does support browsing, though to a limited extent; the limitations of Greenstone's browsing facilities are discussed in Sect. 3. Greenstone is used by numerous organizations worldwide to manage and present collections of documents.

The aim of this work is to present possibilities for adding greater flexibility to Greenstone's browsing facilities in order to provide better support for the information-seeking process. When considering ways to make browsing easier and more effective, a number of potential directions presented themselves: we could provide better facilities for browsing through search results or for browsing through an entire collection, we could present conventional list displays or experiment with 2D or 3D

visualizations, and we could exploit subject metadata (if available) or use learning techniques to extract subject or similarity information.

The first browsing approach presented allows users to specify parameters such as the metadata by which they wish to browse and the maximum number of documents on a page (Sect. 4); this supports browsing list displays of search results. The second approach uses a self-organizing map (SOM) to create a browsable, searchable representation of clusters of documents within a given collection (Sect. 5) – where the SOM uses machine learning to create similarity measures between documents and provides a layered, 2D visualization of the entire collection. The third approach incorporates a subject thesaurus into the search interface to support users in query refinement and in exploration of the subject vocabulary and structure of the document collection (Sect. 6). All three of these enhancements allow the user to easily move between searching and browsing activities – a significant requirement for an effective digital library interface, as this supports common information seeking behaviour [5, 28, 40].

We adopted a formative evaluation approach [31] in developing these three browsing facilities – that is, prospective users were involved in each stage of their design and development, and feedback from these users/testers guided the creation and refinement of the implemented interface. Each browsing facility was assessed for “learnability”: the extent to which a user can get started working with the system without first undergoing training [31]. A high proportion of information seekers are perpetual novices with the (potentially many) information sources that they consult [8], and so a system that can be immediately useful will be more likely to see future use. Assessment of learnability includes evaluation of predictability – the ability of users to predict system reactions [15]. The findings of the user studies, and their

impacts on the design and development of their respective browsing facilities, are summarized with the descriptions of the three browsing interfaces.

The formative evaluation development process is employed to improve the usability of a system or interface, where “usability” refers to the ease with which a person can learn to interact with and then use a system or component of a system [21]. As is argued in Sect. 2, the inclusion of additional browsing features in Greenstone will increase support for a common information-seeking technique. The Technology Acceptance Model (TAM) [13] predicts that user acceptance rates (and hence actual usage) of a system will be higher if the system is perceived to be easy to use (to have a greater degree of usability, brought about through the use of formative evaluation) and perceived to be useful (by adding new, useful features to Greenstone); the intent, then, is to enhance Greenstone so as to encourage a greater user acceptance of it.

Section 2 of this paper discusses the information-seeking process: examining what browsing is, why it is important and how information systems can best support it. Section 3 describes Greenstone’s current browsing capabilities and considers their weak points. Sections 4–6 present the three new browsing facilities, and Sect. 7 draws some conclusions about this work.

2 Human information seeking

The information-seeking process begins with the conception of a need for information, and (if successful) ends with the satisfaction of the information seeker that they have the information they require [26] (though in reality users often “satisfice” [3], or simply give up [34]).

Searching may be most of information *retrieval*, but it is not all of information *seeking* [20]. To be as useful as possible, digital library interfaces should support all stages of the information-seeking process (for example see [7, 26, 34]), not just searching. Despite research having long emphasized that browsing is a fundamental information-seeking activity (Bates described this in 1989 [6]), many systems still do not support it [27, 38]. Moreover, the information-seeking process is not necessarily linear; users generally seek information in an iterative manner, switching back and forth between stages [5, 28, 40], particularly between searching and browsing.

Browsing can be supported by many different facilities, including semantic browsing using such tools as self organizing maps [10] or phrases [36], metadata-based browsing like the Greenstone classifier system [4], and subject categorization (for example the Library of Congress classification scheme). Browsing may be within a document (for example leafing through a book) or between documents (for example wandering the library shelves). Browsing may occur for a number of reasons, including evaluation of an information source, information discovery and clarification of an information prob-

lem [10, 28, 38]. A common definition of browsing is an exploratory information-seeking strategy relying heavily on serendipity and being used to meet an ill-defined information need [6, 10, 33].

One way conventional libraries support browsing is through subject classification of documents. However, physical libraries cannot rearrange the shelves at whim to meet the needs of the user (say if they wanted to browse by author and changed suddenly to title). Electronic information systems (such as digital libraries) have the opportunity to “rearrange the shelves”.

For a system to support browsing effectively, it must be flexible so as to allow the user to modify their information need and information-seeking strategy at will. It should support browsing for any number of reasons, including those mentioned above. For optimum information-seeking effectiveness, interleaving of browsing and searching should ideally be simple [12, 22].

Browsing is easily shown to be a vital part of the information-seeking process and very effective when combined with searching. Information systems need to recognize this importance and support browsing in ways that will allow users to become effective information seekers.

3 Greenstone and browsing

Greenstone is a complete digital library management system, handling everything from collection building to collection presentation via a Web browser. It facilitates full-text and metadata searching as well as various kinds of browsing [1, 4]. Greenstone is designed to allow collections to be built fully automatically (i.e. to not require the manual processing of source documents) and to be served by inexpensive machines over a slow Internet connection [43]. Greenstone is largely stateless, not keeping information about what users do from one action to the next, so as to help reduce server load. Section 3.1 discusses the current browsing facilities available in Greenstone, and Sect. 3.2 explains why these facilities inadequately meet information seekers’ needs.

3.1 Greenstone’s current browsing system

Greenstone’s current browsing system is known as the “classifier” system: documents are placed into classes at collection build time according to their metadata, and browsing structures are pre-built ready for loading. Greenstone supports a number of different types of classifier, each suited to a specific kind of metadata. Each classifier displays information in its own way.

There are five main types of classifier currently implemented in Greenstone: the list, the alphabetic classifier, the hierarchic classifier, the date classifier [4] and the phrase-based classifier Phind [36].

The list classifier is the simplest of the classifiers; it merely sorts metadata alphabetically and presents documents in a single long list (Fig. 1a).

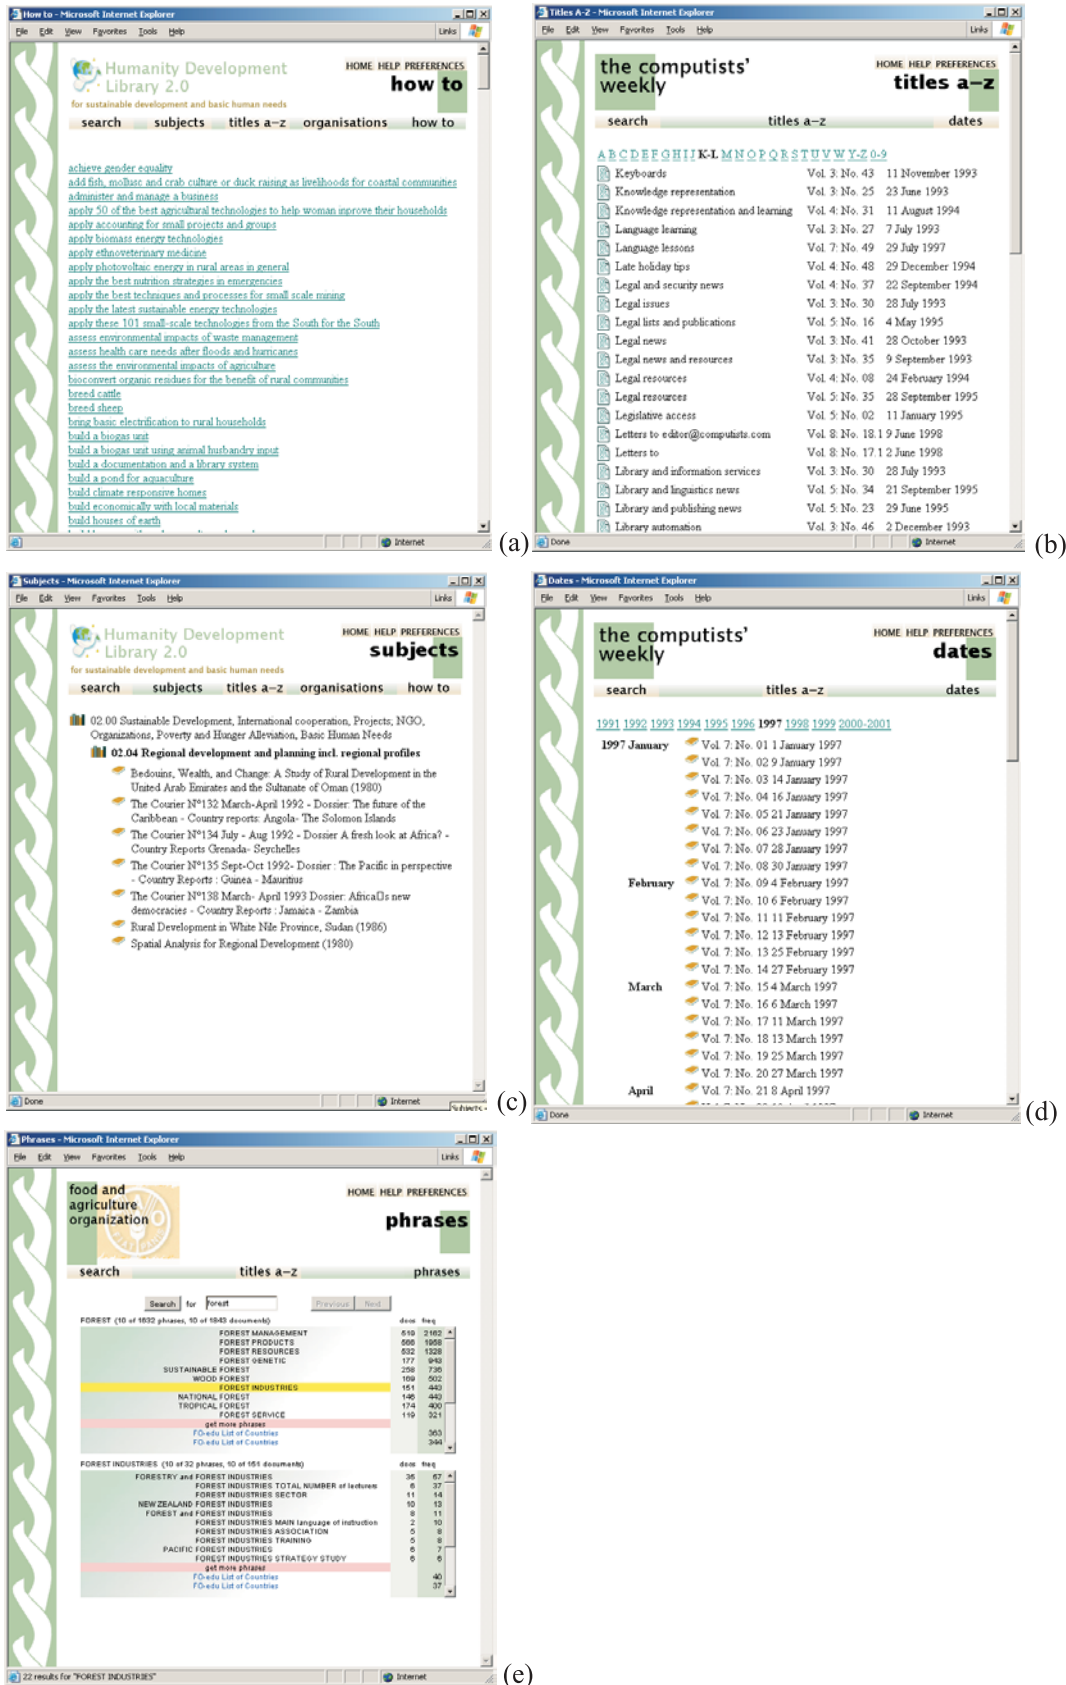


Fig. 1a–e. The Greenstone classifiers. **a** A list classifier of “How to” metadata. **b** An alphabetic classifier, viewed by title. The section being viewed is “K–L”. **c** A hierarchic classifier. The classification being viewed “02.04” is two levels deep. **d** A date classifier. Note the months down the side of the page. **e** The “Phind” classifier. The word “forest” is being browsed

The alphabetic classifier also sorts documents alphabetically, but the document list is then divided into classes according to initial letter, and the classes are displayed across the top of the page (Fig. 1b). If the classes are smaller than a pre-set size, the classifier will merge them (e.g. “K–L” in Fig. 1b). There is no upper limit on the number of documents in a class.

The hierarchic classifier deals with numerical hierarchies – documents are assigned a number indicating their position in the hierarchy (much like the Dewey decimal system), and the user views the hierarchies by progressive drill-down clicking (Fig. 1c).

The date classifier is very much like the alphabetic classifier, though it uses the year as a basic unit (as opposed to initial letter) and displays month information down the side of the hierarchy (Fig. 1d).

The Phind classifier is not based on traditional metadata. Instead, it creates an index of phrases when the collection is built and allows the user to browse by entering a single word or phrase and drilling down through phrases to documents (Fig. 1e).

3.2 Problems with browsing using the classifiers

The classifier system in Greenstone does not support users as well as it might. The failings are in two major areas – the fact that users cannot combine searching and browsing, and in the rigidity of the system.

As discussed in Sect. 2, users locate information most effectively when they can switch easily between searching and browsing. Search results in Greenstone are currently always displayed in lists, and if a search is not ranked, these lists are unsorted. Classifiers present all the documents in a collection that have the classification metadata; there is no way to search a classifier. Thus the cognitive cost of switching between searching and browsing is high, reducing information-seeking effectiveness.

The rigidity of the classifier system is built-in – each classifier uses static, pre-built browsing structures, thus allowing collections to be presented only in one pre-determined manner without input from the user. To illustrate how this can become a problem, imagine a collection with a number of distinct documents with the same title and different authors. Users cannot specify that they would also like to see the author metadata when browsing, much less insist upon the documents being sorted by author. Another example of how the rigidity of the classification system is detrimental to the information seeker’s experience is the size of the groups displayed. Users are better able to navigate and evaluate options if they do not have to scroll [32]; yet in medium-sized collections (say 1000 documents) users may have to scroll through three screens of a single classification, and there is no way for users to specify the largest number of documents they wish to see on a page.

The Phind classifier solves the rigidity problem, but allows browsing of only a single kind of metadata

(phrases), and still does not allow collections to be filtered by search terms – and thus does not entirely solve the browsing problem.

Greenstone supports browsing in a limited way: non-searchable, static metadata classifiers. While this approach goes some way towards supporting browsing, it hinders users in their information seeking by not allowing them the flexibility necessary for truly effective information seeking. Moreover, Greenstone has strong goals related to usability, utility and simplicity of collection creation. The work described in Sects. 4–6 is an attempt to overcome the failings in Greenstone while still taking into account its goals of simple collection provision on inexpensive hardware.

4 A new search-based browsing system for Greenstone

This first browsing enhancement to Greenstone focuses on an inter-document metadata-based browsing system. This approach was chosen because Greenstone is about presenting collections of documents (rather than single documents), and Greenstone already has an effective semantic browsing system in Phind [36]. The user capabilities the new system was to support were defined with the failings of the current system and the research on human information-seeking behaviour in mind. These capabilities are as follows: users must be able to combine searching and browsing, users must be able to choose the metadata by which they browse, users should be able to browse by more than one kind of metadata at a time, and users should be able to restrict the amount of information on any one screen. The guiding principle is to give the user the richest possible browsing experience. This browsing system is described in Sect. 4.1, and its evaluation is discussed in Sect. 4.2.

4.1 The search-based browsing system and its implementation

This new browsing scheme is designed to provide a rich browsing experience without being too taxing on the user – allowing users flexibility in specifying how they wish to browse, while providing a simple interface with sensible defaults. This also involved enhancing the existing browsing system to handle alphanumeric and date metadata.

One of the major advantages of this system over the existing classifier system is that searching and browsing can be combined easily. The search offered by this interface is functionally identical to the ordinary Greenstone search, but results are presented in a browsing structure defined by the user. To avoid the loss of the useful ranking information provided by Greenstone’s underlying search technology MG [44], where the search is an “any words” search, rank information is displayed next to the document metadata, similar to search engine results (Fig. 2b). If the user does not enter any search

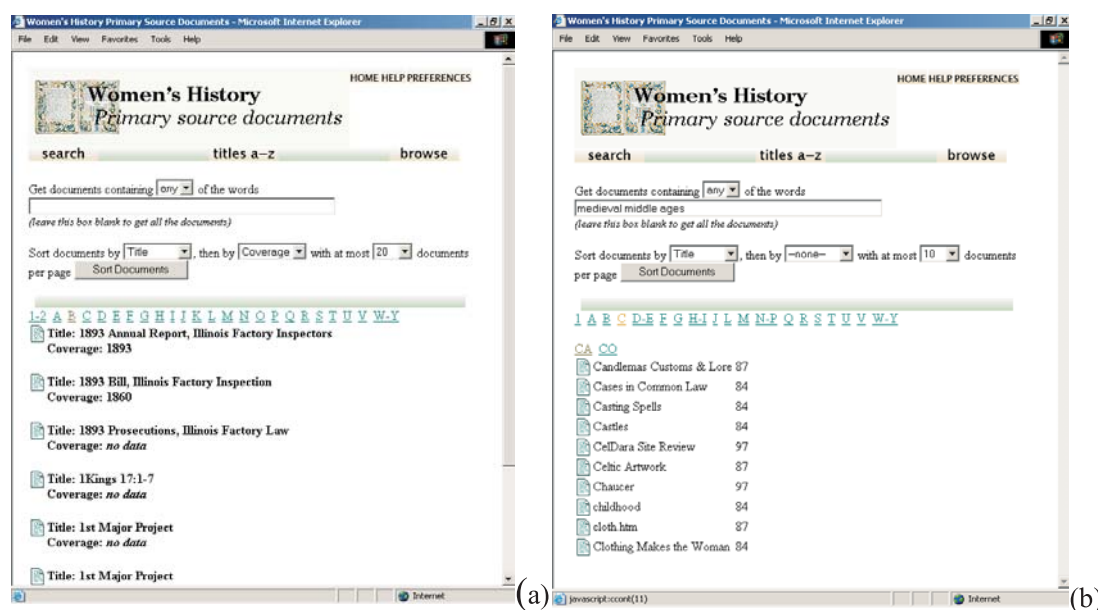


Fig. 2a,b. The new search-based browsing system. **a** Browsing of a whole collection by two pieces of metadata. **b** Browsing of ranked search results by title metadata. Note the ranking information on the right and the second level of the hierarchy

terms, then the user can browse the whole collection (Fig. 2a).

The mechanism for specifying how documents are to be browsed must allow great flexibility, but it also must be simple enough to use without training. To that end, the user is presented with familiar Web-based controls and simple language to determine their browsing preferences. They may browse one or two kinds of metadata at a time (the lists of metadata available for browsing are requested from the collection and inserted into the interface when the browse page is displayed), and they may specify how many documents they wish to see on a page.

Because Greenstone is stateless, and combining all this information to form a browsing structure is computationally expensive, the browsing structures are created only once, and the classes not being currently viewed are hidden in the page using dynamic HTML and JavaScript. This means that the user will get an instantaneous response when switching between the classes in a browsing structure.

Browsing more than one kind of metadata at a time allows the user to view distinguishing metadata where the primary browsing metadata values occur more than once (for example many books with the same author but different titles, when browsing first by author). It also allows the user to sort the duplicates by the second piece of metadata (so sorting the books by author, and then sorting the books with the same author by title). Both pieces of metadata are displayed for each document, even where documents may only have one of the two pieces (Fig. 2a). Documents without the first piece of classification metadata are slotted into the browsing structure under the label “no metadata available”.

The system accommodates the “number-of-documents-per-page” option by creating a two-level hierarchy, with basic divisions (for example initial letter in the first level of the hierarchy, and the second level of the hierarchy divided up so as to provide the required number of documents in each class). The classes at this level are labelled such that they are distinct from neighbouring classes (for example if the first document in one class is called “teak chests” and the first document in the next class is called “teas of the world”, the classes will be labelled “teak-” and “teas-” respectively). See Fig. 2b for an example of such a browsing situation.

Users rarely change the defaults on information-seeking interfaces; therefore, while this interface offers a lot of flexibility, it must also have sensible defaults [24]. By default, the entire collection is displayed for browsing by title metadata; this default is chosen with a view to giving the user a good overview of the collection. The default second piece of metadata by which to browse is determined by whatever metadata the collection has – it is hard to tell automatically what will be useful for any given collection, so the default second piece of metadata is the first detected piece of metadata that is not the title. The default number of documents per page is 20; this is approximately one screen full, so as to avoid wasted screen real estate, but also to lessen the need for scrolling. Searching defaults to “any” words, as this is less likely to give a “no match” result and therefore less likely to frustrate the user [24].

The new browsing system has been designed to offer flexibility in a very simple manner. The major advantages it has over Greenstone’s existing classifier system are the combination of searching and browsing and the ability to

interactively change browsing structures to meet changing information needs.

4.2 Evaluation

There were three main components to the evaluation of this system: a technical evaluation (Sect. 4.2.1), a user study of the way users would like to browse (Sect. 4.2.2) and a user study testing the predictability of the system (Sect. 4.2.3).

4.2.1 Technical evaluation

Systems must be technically sound to be useful. There are two aspects of the new browsing system that can be meaningfully evaluated: scalability and time constraints.

The new browsing system attempts to address the scalability issues faced by the old system (i.e. browsing lists potentially growing very long in medium large collections) by introducing a second level to the browsing hierarchy. Consider a collection with 26 000 documents, with the initial words of titles evenly distributed through the alphabet; the title browsing interface of the old system would display 1000 documents in a long list, for each letter of the alphabet. The new system would divide these classes of 1000 documents up into smaller classes of documents, containing, say, 50 documents each (determined by the “maximum number of documents per page” setting on the interface). This means there will be 20 subclasses across the top of the page under the top-level classes, which is still reasonably usable. Of course, it is possible with this system to have so many documents that it also becomes unusable, but this upper bound on a manageable collection size is much larger than in the old system.

The time constraints on the new system are more worrying. Users hate waiting for Web pages to load [24, 33], so load time has a large impact on the usability of a page. Unfortunately, because the new system produces pages that contain entire browsing structures, the pages are very large. The browsing interface itself is 6.78 KB and each document is 0.04 KB in the browsing structure. This means that over a 56 kbps modem the interface will take 1 s to load, and each document in the browsing structure will add about 0.05 s; with a large collection this adds up very quickly. A collection of 1073 documents (browsed by title and coverage) was shown to take 66 s to load over a 56-kbps modem connection, precluding this interface from being used over a low-speed connection. However, for a high-speed connection or a local collection, this time drops to under 1 s, and once the page is loaded the entire browsing hierarchy is available instantaneously. When we compare this to the classifier system, the total load time over a 56-kbps connection would be 88 s for title metadata only. However, this time is in smaller chunks as the user loads each part of the hierarchy, and thus the wait time is more palatable to the user (each individual page would take about 3 s to load over a 56-kbps connection).

A transaction log analysis of 42 collections in the New Zealand Digital Library [2] from 21 June to 19 December 2001 shows that approximately 9.5% of all actions are browsing with the classifier system and that 37% of the time when users looked at one part of a classifier (say the “A” section of a title classifier) their next action was to look at another part of the same classifier (say the “B” section). This has an associated time cost under the old system, but under the new system it is instantaneous.

4.2.2 How users would like to browse

The first user study was designed to examine how, given free rein, users would like to browse a collection of documents in an online information system.

The first study selected ten participants from the summer population of the Department of Computer Science at the University of Waikato. Participants included tutors, lecturing staff, and graduate and undergraduate students. All the participants had some research experience (the undergraduate students were employed in summer research positions at the time of the study), though not necessarily in the digital library or information-seeking fields.

Each of these users completed a questionnaire about on- and offline browsing and their experience with it. All of the users had used online search engines, and most felt that they “usually” found what they were looking for with these tools. Only two of the participants had more than basic experience with online browsing tools, and many commented that they “didn’t like them” (though one user who had a lot of experience with them liked them a lot and commented they were good for “filtering out crap sites”).

After completing this questionnaire, users were given a set of 30 index cards, each of which represented a document randomly chosen from the Women’s History Collection in the New Zealand Digital Library. These index cards each showed four types of metadata, as recorded in the online collection (see Fig. 3 below). Participants were asked to arrange the documents the way they would wish to browse them if they were looking for information on medieval women.

When participants had arranged the cards to their satisfaction, the arrangement of the cards was manually recorded, and participants were interviewed about their

```

Title: Shirtwaist Strikers in New York
Author: no data available
Coverage: 1912
URL: http://www.binghampton.edu/whist/shw/doc12.htm

```

Fig. 3. An example of an index card used in studies of the search-based browsing system

browsing structure (that is, a retrospective verbal protocol was applied [41]) and their expectations of an online information-seeking system and what it would allow them to do.

Of ten study participants, six created a purely metadata-based browsing structure, two created purely semantic browsing structures, and the remaining two created structures that were primarily metadata driven but had some semantic elements (semantic in this context is used to mean “based on meaning inherent in the metadata, rather than on the metadata itself”).

The semantic structures were not reproducible by machine; one merely sorted the cards into two piles of “useful” documents and “not useful” documents, and the other was based on over-broad and vague categories such as “law and politics” and “sociology”.

The metadata-driven structures were based variously on coverage metadata and title metadata. Only one structure was purely alphabetical based on title; the remainder used a combination of the two kinds of metadata.

When questioned about what they would expect from an online browsing system, three users said that they would expect a search facility as well as browsing support (in keeping with the users’ abundant experience with online search systems and lack of experience with online browsing systems). Three participants commented that “an information system should be able to present more than one view of the documents”, “I would like to be able to sort by different metadata if I wanted to” and “you can’t rearrange the shelves in a library, but I would expect an information system to be able to arrange the documents according to all kinds of metadata”.

In sum, the users predominantly wanted metadata-based browsing systems, wanted to be able to combine search and browsing, and wanted to “rearrange the shelves” to browse by different kinds of metadata. Each of these functionalities is something that search-based browsing offers over and above the conventional Greenstone browsing experience.

4.2.3 Predictability of the search-based browsing system

The second user study was designed to ensure the usability of the search-based browsing system. It has been shown in research that users do not use help systems when using digital libraries, and that if something is tricky to learn, it will be rapidly abandoned [25]. With this in mind, it was important that the search-based browsing system produce results that the users would expect, given the interface.

There were eight participants in this study, selected from the same pool of people as for the first study. The experiment used the same set of 30 index cards described in Sect. 4.2.2. In this experiment, however, users were shown a picture of the interface for the search-based browsing system (Fig. 4) and asked to arrange the documents the way they believed the system would arrange them. The layout they generated was manually recorded, and then the participants were asked to explain what about the interface led them to believe the system would create a similar structure (a retrospective verbal protocol [41]).

Seven of the eight participants sorted the documents much in the way the system would, with a few minor differences (for example sorting punctuation before titles beginning with “a” and commenting that “the system shouldn’t do it this way but it probably does” – it doesn’t). The eighth participant grouped the documents by first letter in the title and then sorted the groups by coverage metadata. When asked what about the interface made him believe documents would be sorted this way, he replied, “This is not how I believe the interface would sort the documents; this is how I would like it to sort the documents – in reality the interface would sort the documents alphabetically”.

All eight participants in the second study were able to identify how the interface would sort the documents, either according to the experiment protocol or verbally afterwards. This implies the search-based browsing system behaves in a way users find readily explicable, and

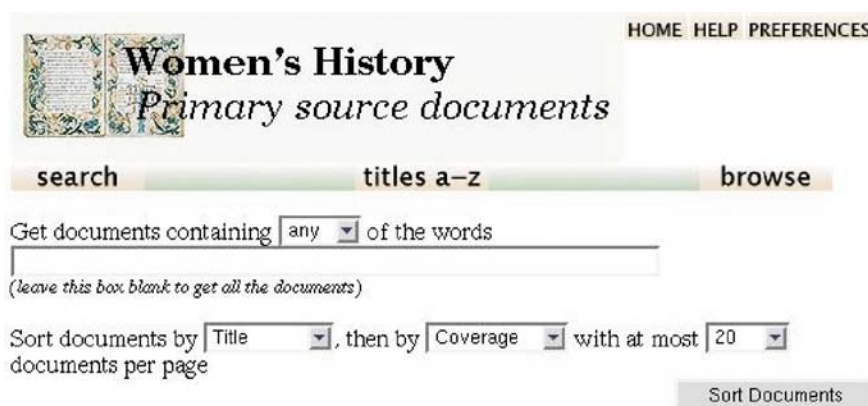


Fig. 4. The interface picture shown to participants in an experiment determining the predictability of the search-based browsing interface

thus that users will not need to consult online help to understand the interface.

Evaluating the new browsing system both technically and with user studies shows that it has only one major flaw: the amount of time a browse page may take to load (and even that is ameliorated by the fact that it then provides better performance than the standard browsing interface 37% of the time). The new system handles large numbers of documents better than the old system, is readily comprehensible, and allows users the flexibility they want in a browsing system.

5 A self-organizing-map-based browsing facility

This section describes a layered, 2D interface for browsing a digital library collection, based on a self-organizing-map (SOM) representation of documents. A SOM is an unsupervised machine learning algorithm used to cluster and visualize high-dimensional data sets [10]. It is particularly attractive for inclusion in digital library systems lacking human-assigned subject classification metadata because it creates an automated grouping of “like” documents and presents this clustering in a hierarchic, map representation. At present, it is often difficult to quickly get an overview of a digital library information space or to visualize the relationships between documents. The map representation is intended to support browsing with an easily understandable interface that loads quickly on standard Web browsers. Of course, a user should not be faced with an either/or choice between keyword searching and browsing; people often employ both strategies when engaged in information seeking, for example in the common “berry-picking” behaviour [6]. The implementation of the SOM-enhanced Greenstone interface allows the user to move easily between keyword searching and SOM-based browsing.

5.1 Implementation and interface

As noted in Sect. 3, the Greenstone user interface has significant limitations: searching and browsing are conducted in separate screens, and it is difficult to move between the two activities. Additionally, subject browsing is based on metadata supplied by the collection creator/maintainer; there is no built-in mechanism for automatically clustering “like” documents together, and so subject browsing is not supported in the absence of subject metadata.

The SOMLib digital library system [14] uses the GHSOM algorithm [37] to create a hierarchical, content-based organization of document collections to aid browsing. SOMLib provides two visualizations of this information: a text-based two-dimensional grid, and a graphics-based presentation of documents as “books” on shelves. SOMLib offers support for building browsing interfaces but has no effective search facilities.

The browsing facility described in this section brings the two technologies, Greenstone and SOMLib, together. SOMLib is used to cluster the documents. The maps for a collection are then passed, together with the documents themselves, to Greenstone for collection indexing and search interface construction. The maps are linked to the “browse” button on the Greenstone toolbar by replacing the Greenstone browsing functions with instructions to display the appropriate map file. A Greenstone “plugin” automatically creates document metadata describing which lower-level map a particular document is in – and so each document is linked to a map that contains all the other documents that it has been clustered with.

Figure 5 presents a top-level map for a sample collection, and Fig. 6 illustrates the display of a lower-level map. Context information is given at the top of each map, describing which level of the hierarchy is represented (the “layer”) and how many documents are reachable “below” this map. Navigation up and down the layers is accomplished through either the “Up/Down one level” links or by the “View all documents” link.

Each cell in the 2D grid represents either a single document or a group of documents. The contents of each cell are described by a list of automatically extracted keywords that are common to the documents clustered together in the cell. The most heavily weighted of these keywords is displayed bolded at the top of the cell; for greater readability, the remaining keywords are listed in alphabetic order.

The map can be reached through the keyword search facility of Greenstone; search result lists include a “View all related documents” link with each document. Clicking on this link leads to the map containing that document, with the original document highlighted within its grid. The user now has the opportunity to continue browsing



Fig. 5. Portion of a top-level collection map

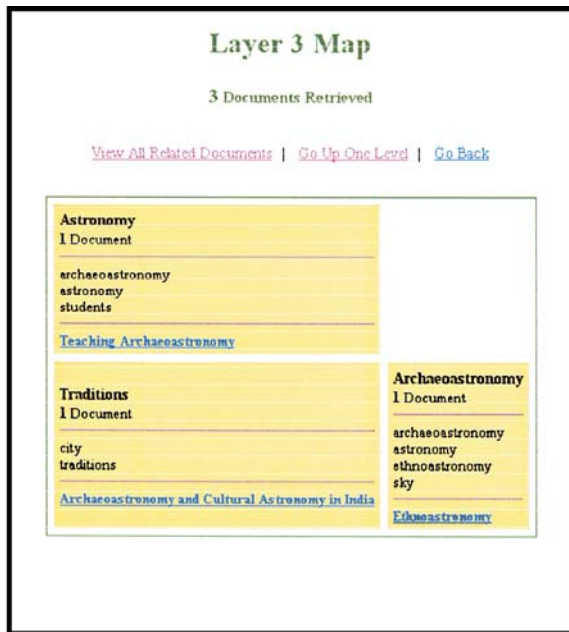


Fig. 6. A bottom layer map

using the map interface, go back to the linear search list, or view selected document(s) through the grid. The presence of this option is key to interleaving browsing and searching.

The SOM clustering is visualized using an enhanced version of SOMLib's grid interface rather than the graphical bookshelf display. The grid interface was appealing because of its greater efficiency – the user is not faced with the annoying time lags common with graphics-intensive visualizations. The text-based interface also has the advantage of using only HTML in its display, so that the user is not required to download plugins or applets. Additionally, there is no experimental evidence to show that a detailed 3D presentation improves the effectiveness of browsing – and given that the 3D bookshelf view is more difficult for a reader to scan than the primarily text-based grid display, the 3D display is likely to make browsing less, rather than more, efficient.

The original SOMLib grid was tailored to the task of browsing document collections to improve its readability and aesthetic appeal: the maps are now centred on the page, cell borders and spacing have been altered to provide greater definition between cells, colours are now used consistently and indicate depth within the map, keywords are sorted alphabetically, SOMLib parameter values have been removed from the display and larger font sizes and bolding are used for drawing the eye to cell labels. Navigation has been improved by providing an option to view all documents mapped to a cell without forcing the user to navigate down to the lowest levels of the hierarchy and by adding a count of the total number of documents accessible through each unit below the unit's label (to give a sense of the size of the cluster yet to be explored).

Note that this map is of the entire collection contents and not of the results of a search. Creating a SOM of search results would take varying – and possibly unacceptably long – amounts of time depending on the number of hits returned by a search, and users dislike variable response times [39]. Users of search engines and digital libraries also rarely look beyond the first page or two of search results [24] and prefer ranked listings of results; there is scant evidence that visualizations of search result distributions are a desired feature in a digital library. For these reasons, we focused on providing a visualization of the collection as a whole. The ability to move from a search result to its associated SOM cell supports serendipity in browsing, as the SOM brings to the user's attention documents that are similar to the initial search hit but that may not share the search terms entered by the user.

5.2 Usability evaluation

Two user studies were conducted: the first to investigate whether the SOM map display is likely to fit comfortably with the way potential users would wish to organize documents (Sect. 5.2.1) and the second to provide usability feedback on the SOM interface (Sect. 5.2.2).

5.2.1 Acceptability of the SOM visualization

The first user study was a card-sorting exercise similar to the study described in Sect. 4.2.2 – participants sorted document surrogates into clusters. Here, however, the focus was on spatial layout of the groups and whether the SOM collection visualization reflects the way people group documents spatially.

Twenty participants were recruited from the local computer science students and faculty. A set of 20 laminated cards was prepared, each containing a summary of a science article downloaded from the Google directory (<http://directory.google.com>). Each summary included the article's title, a set of keywords taken from the document and an introductory (summative) para-

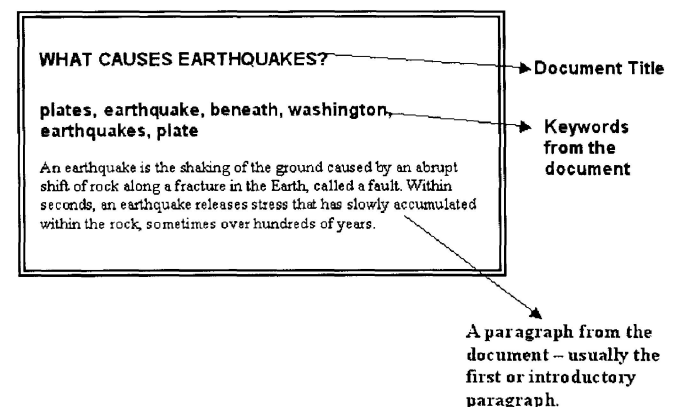


Fig. 7. A sample card used in the SOM card-sorting study

graph from the document (Fig. 7). The cards were numbered randomly, and the stack of cards received by each participant was ordered in exactly the same way. The card number was used at the end of each session to record where the participant placed each of the cards.

Each participant was requested to imagine that they had just finished gathering documents of personal interest from different fields of science, represented by the 20 cards, and now they want to organize the documents so as to make them easy to work with when they are handled next. Participants spent half an hour, on average, grouping the cards.

Upon completion of the task, a photo was taken of the final arrangement chosen for the cards. Each participant was then asked to explain the document grouping that they had created (a retrospective verbal protocol [41]).

All 20 participants placed the cards in “the most convenient place to reach” (participant Q). Sixteen of the participants chose to place their cards in horizontally arranged groups, even when not constrained by the area in which they were to place the cards. Two of the participants (B and N) were constrained by the area in which they were performing the task and so opted for a more vertical layout. The remaining two participants (I and S) preferred a vertical layout for the cards.

All participants reported that they had attempted to group the cards “by major topic” (participant J), where this subject was self-perceived (sample pictures of participants groupings are shown in Fig. 8). Fourteen participants grouped the cards with the “more general [card] on top, then move to more depth about the topic” (participant C), starting with “big things, then it gets smaller” (participant K). That is, they liked going from a “larger perspective then down to the more specific” (participant M). Participant B remarked that they preferred “putting the documents in the order that I would use them”, but that they had “no sub-sorting unless the groups are quite big”. Size of group was a consideration in

constructing the clusters: participant F stated that they had only “broken down [groups] into sub-groups if [they] got more than 5 cards” in a group, and two participants stated that they “don’t like having single items in a group (participant B) because they “don’t like seeing one thing by itself” (participant F).

Eight participants chose a spatial layout that attempted to show a relationship between different groups by placing “similar kind[s] of things together” (participant S) and using “the gaps in between things” to “represent relativity” (participant H). Participant Q corroborated that sometimes, wherever possible, “the spaces show how related the documents or groups are” and that “related documents in a topic are arranged vertically”. Participant O attempted to “establish indirect links between items” because “forming relationships may help remembering information”. The remaining participants preferred to have their cards “plonked randomly” (participant N) wherever they found room for them. Once the cards were arranged, the minimal amount of information that each participant wanted to see in each card was the title, or if only the top card of a group could be seen, then all of the information on that card.

In summary, the layout employed by all of the participants consisted of discrete clusters, and most of the participants arranged these clusters in a predominantly horizontal display. Some of the participants attempted to use the space between the various clusters to represent relationships between the clusters, where the larger the gap between two clusters, the less the clusters were seen to be related. Most of the participants also tried to create an ordering of the cards within each cluster. All of the participants placed the groups within easy sight and reach. None of the participants employed a linear layout, and only one participant used an alphabetic ordering as the predominant arrangement.

Consider the way results of the SOM are visualized in a two-dimensional, map-like layout. Related documents

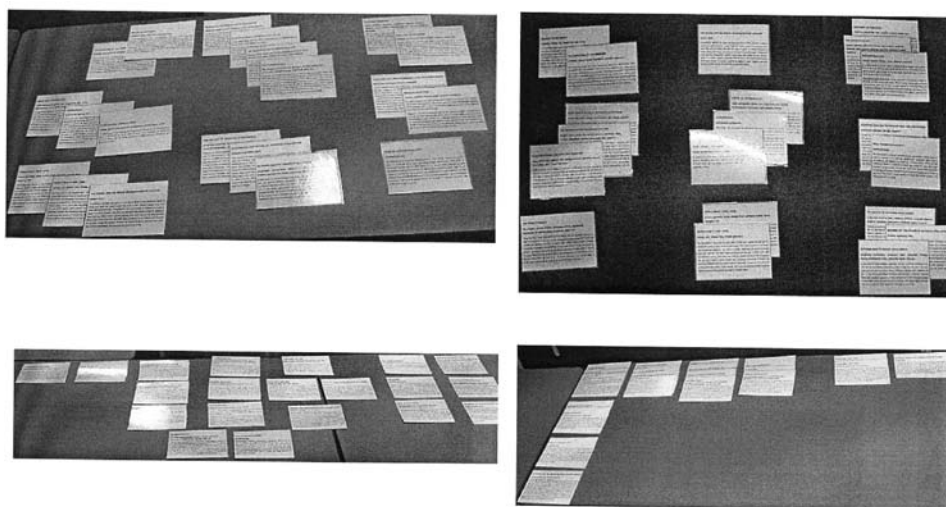


Fig. 8. Sample photos of final card layouts for participants in the SOM card-sorting study

are grouped together, and related groups are presented in adjacent cells. The participants organized their cards similarly. Most participants employed a predominantly horizontal spatial layout; the SOM display parameters were set to expand more horizontally than vertically to accommodate this observation. Most of the participants created a hierarchy within groups either by stacking the cards in a general-to-specific structure or by creating sub-groups ordered in a similar general-to-specific manner. This hierarchical structure is accommodated by the SOM layers. Some of the participants commented that they do not like traversing through more than approximately three levels in a hierarchy. Taking this into account, the number of layers in the SOM display has been limited to three.

5.2.2 Usability study of SOM display

A usability study was conducted that focused on how new users interpret the map display. Given that digital library users generally resist reading interface instructions or help files, users must be able to quickly deduce the meaning of this relatively novel interface. As recommended by [30], this was a small-scale usability study, with five participants. (Using more than five participants in an initial usability study tends to be wasteful, as the same problems are identified over and over again.) The participants were chosen from the pool of participants who took part in the study described in Sect. 5.2.1. The participants were each given two tasks. In the first, they were asked to examine the SOM display and to give their opinion as to the functions and meaning of the various screen elements and to state what they believed would happen if they clicked on the various links and buttons on the page. Written notes were taken of their responses (a concurrent verbal protocol [41]). In the second task, participants were asked to use the SOM-based browsing interface to find a set of documents and to “think aloud” as they performed the task (concurrent verbal protocol [41]). After this second task was completed, the participants were de-briefed about their opinions of the SOM display.

The participants pointed out shortcomings in the labels and keywords generated by the SOM: participants found them to be very “unhelpful and confusing” (participant iii) because “they are the same in different units” (participant iv) and did not “match the documents they referenced” (participant i). All participants were confused when the SOM generated the same label for different units. Participant i added that “the labels are not descriptive enough”, and participant iii concurred that they want the labels to be an “accurate phrase that summarizes the document”. These comments point to an open problem in SOM research: the difficulty of automatically generating human-comprehensible labels for SOM document clusters [10].

The participants were not able to establish whether there was a relationship between units that were placed

in close proximity to each other – mainly because the labels of neighbouring cells did not make this relationship clear. The participants therefore decided that the units were placed randomly because units with “similar labels are not placed together” (participant i).

All participants could understand what the different coloured units meant; participant i stated that they “understood the set-up very quickly because [the interface] has organized everything for me”. Apart from dissatisfaction with the keyword labels, all of the participants found the other parts of the map interface easy to understand. Some expressed concern that they “did not immediately understand what ‘layer’ or ‘level’ meant” (participant v), but it was “easy to make out after exploring the links” (participant iv).

All the participants commented that they wanted more information about where they were in the hierarchy to be indicated within each layer of the map; this is an opportunity for future development of the interface. All of the participants liked the fact that they were allowed to “view all documents” that can be accessed through each unit with one click. They suggested that this option should be included with the layer 1 map as well.

The results of the usability study indicate that, though participants had difficulty interpreting the automatically generated keyword labels, the SOM-based interface provided them with a “good overview” (participant iv) of the documents in the collection. The participants also felt that the interface was “easy to follow” (participant v) “once you played around with it a bit” (participant i). While the relationships between clusters were not generally understood – that is that neighbouring clusters are more similar than distant clusters – it was clear that documents within a cluster were in some way similar. All participants liked the fact that the interface displayed quickly and that it allowed them to get a quick overview of the topics represented.

6 Adding a thesaurus to Greenstone searching/browsing

One of the issues that users face when creating and refining queries is the “vocabulary problem”, which occurs when the terms selected by users do not match the terms used in the document collection or its metadata [8, 18]. One useful tool for query construction is a subject thesaurus describing the vocabulary of the collection’s domain and the relationships (typically *broader than BT*, *narrower than NT* and *related to RT*) among the thesaurus terms. In practice, however, thesauri have been little used in searching; subject thesauri are generally not readily available in digital format or, if they are available online, are not integrated into the searching/browsing facilities of the document collection.

6.1 Implementation and interface

A prototype system was developed integrating the AGROVOC [17] subject thesaurus with a Greenstone searching interface. AGROVOC is a multilingual, controlled vocabulary thesaurus developed by the FAO to cover concepts and terminology in agriculture, forestry, fisheries, food, and related domains. It was chosen for experimentation because it is of a size large enough to be challenging for the user (over 16 500 descriptors) and its structure is the common broader/narrower/related-term organization, so it can be used without sacrificing generality of approach in interface design. Additionally, the researchers had access to both a digital version of AGROVOC and a document collection compiled by the FAO in the subject area of the thesaurus. Note that a generic thesaurus such as WordNet could also be incorporated into Greenstone using the same interface described in this section; the appeal of AGROVOC is that it will allow us in the future to investigate subject thesaurus use in searching.

The Greenstone search interface is augmented by suggesting potentially useful thesaurus terms through the search page (Fig. 9) and by adding a thesaurus tab to support exploration of the thesaurus structure itself (Fig. 10).

In Fig. 9, the user has used a term matched by only one document in the collection and may wish to consider additional terms from the thesaurus – presented in the “thesaurus quicklist”. Ticking a quicklist term adds it to the terms in the search box. The quicklist terms are in a broader/narrower/related-term relationship with the

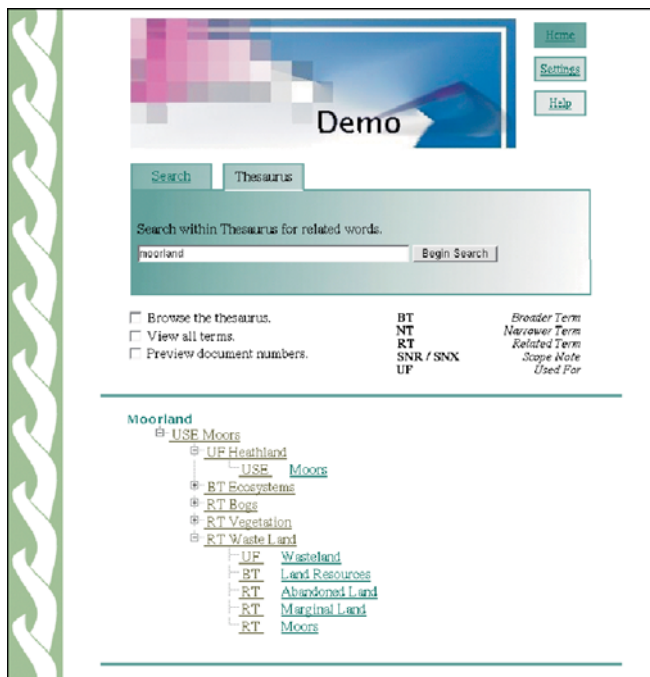


Fig. 9. The thesaurus-enhanced Greenstone search interface

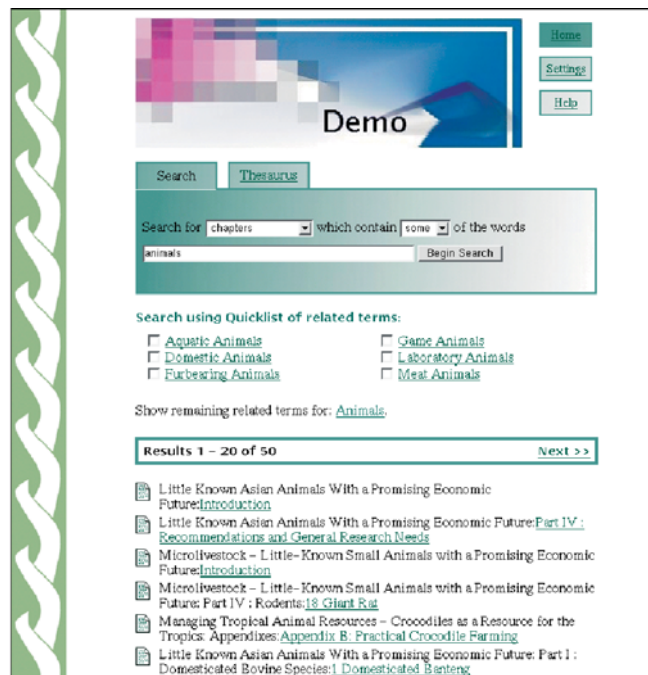


Fig. 10. Searching and browsing the thesaurus

query. If more than one query term is present, then the quicklist contains up to three thesaurus terms for each query term; this limits the total number of thesaurus terms appearing on the search page so as to avoid overwhelming the user with alternatives and to preserve the primacy of the search and hitlist on this page.

Thesaurus terms presented out of context of the thesaurus structure may be less useful to some users than if they were viewed imbedded in the thesaurus. Further, viewing and searching the thesaurus directly can be useful in familiarizing a user with the subject domain of the document collection. To support exploration of the thesaurus, users can browse it directly through the thesaurus tab (Fig. 10). “Child” nodes in the hierarchy can be hidden or expanded in the conventional manner by clicking on the parent term’s +/- icon. The user can quickly focus on a specific portion of the thesaurus through “Search within thesaurus for related terms”.

6.2 Usability evaluation

The thesaurus/search engine interface for Greenstone underwent a total of four usability studies before, after and during its design and implementation.

6.2.1 Establishing guidelines for the initial design

An initial user study provided an indication of preferred interaction style when searching using the thesaurus. The user study evaluated five thesaurus/search engine prototype interfaces representing the range of interface options suggested by a review of the research literature and commercial products. Interfaces I and II provide a single

window to access both the thesaurus and a search engine. In interfaces III, IV and V, the search engine and thesaurus are treated independently in separate windows. Interface IV uses an interactive, graphic representation of the thesaurus, interface V presents the thesaurus as a hierarchy, and interfaces I, II and III use an alphabetic list display of thesaurus terms. In interfaces III, IV and V the user manually searches the thesaurus for terms of interest; the thesaurus is searched automatically in interface II when the user enters a search term, and in interface I the system unobtrusively suggests additional search terms from the thesaurus.

Eight computer science students were recruited as participants for the study. Each participant performed a search on three (randomly selected) interface prototypes and was instructed to “think aloud” while searching (concurrent verbal protocol [41]). Participants’ comments were recorded on video and their screens datalogged. On completion of the session the participants were encouraged to summarize their reactions to the three interfaces (retrospective verbal protocol [41]).

A common theme from the participants was a dislike of slow interfaces (interfaces I and II elicited this reaction most strongly). Slow response time frustrated, distracted and sometimes distressed the participants. When participant 6 was using interface II, for example, it did not seem to respond to his (repeated) clicks. He quickly became agitated, and soon decided that he’d had enough of that interface.

Participants reacted negatively to large amounts of information displayed on a single page, confirming earlier research indicating that users feel uncomfortable when confronted with a “busy” page [19, 42]. Reading and extracting useful information becomes a daunting chore, with one user aptly pointing out that the information he was searching for was approximately 1 cm, on a page with a scrollbar that was 3 mm.

Users do not like switching between independent windows and applications when searching [35]. Participants used the thesaurus less frequently when it was in a separate window and found the multiple independent window interfaces awkward to use. As an example, participant 6 spent as much time rearranging two windows to fit comfortably in the screen as he did searching the collection. One participant confirmed that “you’re more inclined to use it if it’s on the same screen . . .”.

Participants preferred that the thesaurus act semi-automatically (that is, that it suggested search terms) rather than automatically inserting thesaurus terms into the search or forcing the user to manually search the thesaurus for terms of interest. The semi-automated thesaurus is “sort of like when you spell things wrong in Google, it says ‘did you mean?’”; participant 4 noted that “. . . it was good bringing up synonyms whenever I started a new search . . .”.

The visual thesaurus (interface IV) was entertaining – participants spent time playing and having fun with it.

While they enjoyed using it, they felt that in the long term the novelty would wear off and that it would become annoying: “. . . the visual aspect helped a little, but the proliferation of concepts was too much . . .”; “the fact that the thesaurus was in a different window and it was just so darn cool distracted me quite a bit . . .”.

Participants wished to see thesaurus term relationship indications (that is BT, NT, RT) when browsing the thesaurus. They appreciated that they could take in the relationship between a series of terms at a glance.

The study did not provide any indication of whether users felt comfortable opening documents in the same or in new windows. As the goal of this research was to implement a thesaurus searching/browsing interface in Greenstone, the Greenstone default was followed by opening links in the same window.

6.2.2 Testing initial design concepts

Having elicited an initial set of interface guidelines, the next stage in the development process was iterative redevelopment of the interface design [31]. The results from the user study of Sect. 6.2.2 were used to design four visually distinct paper prototypes. These paper prototypes were evaluated by three local experts in HCI and digital libraries, in sessions totalling 5 h. Expert evaluation has economic advantages over other methods for detecting usability problems (for example full-scale usability tests). Nielsen points out that “single” experts are likely to find over 40% of usability problems, but “double experts” – those with “expertise in both usability in general and the kind of interface being evaluated” – are likely to find a higher proportion of the problems [31]. On the basis of this analysis one of the paper prototypes was selected for further development in a low-fidelity, screen-based prototype; this interface prototype was subjected to further evaluation by the three local HCI/DL experts, and on their suggestions minor alterations were made to the interface’s appearance and functionality.

At this point, six students in computer science were recruited to participate in a usability study of the modified, low-fidelity prototype. Participants were asked to “think aloud” as they performed a walk-through of the prototype (concurrent verbal protocol [41]); after completing this walk-through they were queried about the system’s expected reaction to their actions. Written notes were taken of the think-aloud sessions and the de-briefings (retrospective verbal protocol [41]).

Participants recognized the “search” portion of the interface and (correctly) felt confident that it would return a list of documents. The high level of familiarity they felt seemed to give participants a confidence boost when it came to dealing with novel search elements such as the list of alternative search terms suggested by the thesaurus. This assurance vanished when the participants encountered the thesaurus portion of the interface. When queried about the thesaurus, participants gave tentative

guesses involving synonyms, antonyms and related documents. This uncertainty is likely due to a lack of familiarity with subject thesauri.

One participant brought up the issue of ambiguous-term entry (that is entering unrelated terms such as “pig” and “speaker” in a query). Unrelated-term entry is akin to entering ambiguous terms [45]. The brute force method for dealing with ambiguous-term entry would be to simply display a list of all results found for all terms entered, but this runs the risk of information overload [42]. Alternatively, the system could display the results of ambiguous entry in “categories”. User studies performed by Dumais et al. [16] show that category interfaces are more effective than ranked list displays; this approach was adopted in the refined prototype design.

6.2.3 Expert evaluation of the implemented prototype

The design that emerged from the low-fidelity testing of Sect. 6.2.2 was implemented, and an expert evaluation was conducted to discover the usability problems of its interface. Two local faculty members with extensive knowledge of HCI and experience in the design and evaluation of searching interfaces performed the evaluation. Written notes were taken of their analysis. It was noted by both experts that the thesaurus can return a substantial number of terms, possibly causing information overload. Two solutions they proposed included the elimination of duplicate terms and terms that are not present in any documents in the collection, asking, “... is it useful to have them?”

Interestingly, during the course of the expert evaluations, an expert queried whether or not the thesaurus hierarchy should be search oriented or not, whether searching in the thesaurus should help users find documents (by showing document lists, for example) or concentrate on showing relationships between terms. How the thesaurus was going to be used was important, but he felt that it was more important to focus upon document-oriented users (i.e. those who use the system to help them improve their search and find useful documents).

6.2.4 Usability evaluation of the refined browser

The final user study conducted investigated the predictability and learnability of the implemented system, as modified from suggestions raised by the expert evaluation (Sect. 6.2.3). It was also important to see how novice users (i.e. those who do not necessarily use a thesaurus) reacted to the presence of the thesaurus descriptor key (“BT”, “NT”, etc.). Five local computer science students were recruited as participants. Participants viewed (but did not interact with) the thesaurus searching/browsing screens and were then asked to predict the system’s response to potential user actions (i.e. “What do you think will happen if you click there?”). Subjects were then asked to

interact with the system while “thinking aloud” (concurrent verbal protocol [41]). Written notes were taken of each session.

This study indicated a high level of user uncertainty when dealing with selecting quicklist terms – possibly indicating a need for some more immediate documentation regarding this feature. High uncertainty levels when dealing with the thesaurus were present in this study, although users did note that they could easily get a feel for how to use it (“... the quicklist stuff isn’t really obvious, but once you start playing with it, you can quickly see reactions and understand how it works...”).

Overall, the system received positive reactions from both experts and users alike regarding its learnability and predictability. The interface received positive comments about its fast reaction time, something that is vitally important to any Web-based service [23].

7 Conclusions

Three enhancements have been created to the Greenstone software, designed with human information-seeking needs and the failures of the old system in mind. These new systems allow users to combine searching and browsing, in keeping with both the literature on information-seeking behaviour and the user experiments carried out as part of the work done for this investigation. These enhancements allow the user more flexibility in determining the way in which they browse, which is also in keeping with experimental results and information seeking literature: users can browse by more than one type of metadata and select the metadata by which they wish to browse (Sect. 4); in collections lacking subject classification, metadata users can browse clusters of documents grouped automatically into a SOM structure (Sect. 5), and if a subject thesaurus is available for a collection, users can receive support from the thesaurus in query refinement and search and browse the thesaurus itself to familiarize themselves with the subject vocabulary and structure (Sect. 6). These new browsing mechanisms are a vast improvement over the old Greenstone classifier system, making information seeking easier.

References

1. The Greenstone Software (2003) <http://www.greenstone.org> [accessed 6 March 2003]
2. The New Zealand Digital Library (2003) <http://www.nzdl.org> [accessed 6 March 2003]
3. Agosto D (2002) Bounded rationality and satisficing in young people’s Web-Based decision making. *J Am Soc Inf Sci Technol* 53(1):16–27
4. Bainbridge D, McKay D, Witten I (2001) The Greenstone Developer’s Guide. Department of Computer Science, University of Waikato. <http://www.greenstone.org> [accessed 6 March 2003]
5. Baldonado MQ (2000) Wang. A user-centered interface for information exploration in a heterogeneous digital library. *J Am Soc Inf Sci Technol* 51(3):297–310

6. Bates MJ (1989) The design of browsing and berrypicking techniques for the online search interface. *Online Rev* 13(5):407–424
7. Beaulieu M (2000) Interaction in information searching and retrieval. *J Document* 56(4):431–439
8. Borgman C (1996) Why are online catalogs still hard to use? *J Am Soc Inf Sci Technol* 47(7):493–503
9. Chang S-J, Rice RE (1993) Browsing: a multidimensional framework. *Annu Rev Inf Sci Technol* 28:231–276
10. Chen H, Houston AL, Sewell RR, Schatz BR (1998) Internet browsing and searching: user evaluations of category map and concept space techniques. *J Am Soc Inf Sci Technol* 49(7):582–603
11. Crabtree A, Twidale MB, O'Brien J, Nichols DM (1997) Talking in the library: implications for the design of digital libraries. In: *Proc. 2nd ACM international conference on digital libraries*, Philadelphia, pp 221–228
12. Cunningham SJ, Knowles C, Reeves N (2001) An ethnographic study of technical support workers: why we didn't build a tech support digital library. In: *Proc. 1st joint ACM/IEEE-CS conference on digital libraries*, Roanoke, VA, pp 189–198
13. Davis FD (1989) Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q* 13(3):319–340
14. Dittenbach M, Merkl D, Rauber A (2000) Using growing hierarchical self-organizing maps for document classification. In: *Proc. 8th European symposium on artificial neural networks*. Bruges, Belgium, 2000.
<http://www.dice.ucl.ac.be/Proceedings/esann/esannpdf/es2000-39.pdf> [accessed 25 July 2004]
15. Dix A, Finlay J, Abowd G, Beale R (1997) *Human-computer interaction*, 2nd edn. Prentice-Hall, Glasgow, p 163
16. Dumais S, Cutrell E, Chen H (2001) Optimizing search by showing results in context. *CHI 2001*, pp 277–284
17. Food and Agricultural Organization of the United Nations (1998) *AGROVOC: Multilingual agricultural thesaurus*. <http://www.fao.org/agrovoc/>
18. Furnas G, Landauer T, Gomez L, Dumais S (1987) The vocabulary problem in human-system communication. *Commun ACM* 30(11):964–971
19. Hancock-Beaulieu M (1993) A comparative transaction log analysis of browsing and search formulation in online catalogues. *Program* 27(3):269–280
20. Henninger S, Belkin N (1995) Interface issues and interaction strategies for information retrieval systems. In: *CHI Companion '95*, Denver, CO, pp 401–402
21. Institute of Electrical and Electronics Engineers (1990) *IEEE standard computer dictionary: a compilation of IEEE standard computer glossaries*. IEEE Press, New York
22. Jacso P (1999) Savvy searching starts with browsing. *Online CD-ROM Rev* 23(3):169–172
23. Johnson J (2000) *GUI bloopers: don'ts and do's for software developers and Web designers*. Morgan Kaufmann, San Francisco
24. Jones S, Cunningham SJ, McNab R (1998) An analysis of usage of a digital library. In: *Proc. 2nd European conference on research and advanced technology for digital libraries*, Heraklion, Greece, pp 261–277
25. Knepshield P, Savage A, Belkin N (1999) Interaction in information retrieval: trends over time. *J Am Soc Inf Sci Technol* 50(12):1067–1082
26. Kuhlthau C (1991) Inside the search process: information seeking from the user's perspective. *J Am Soc Inf Sci Technol* 42(5):361–371
27. Lawrence S, Lee GC (1998) Context and page analysis for improved Web search. *IEEE Internet Comput* 2(4):38–46
28. Lazar J, Bessiere K, Ceaparu I, Shneiderman B (2003) Help! I'm lost: user frustration in Web navigation. *IT Soc* 3(1):18–26
29. Marchionini G (1995) *Information seeking in electronic environments*. Cambridge University Press, New York
30. Nielsen J, Landauer TK (1993) A mathematical model of the finding of usability problems. In: *Proc. ACM INTERCHI'93*, Amsterdam, pp 206–213
31. Nielsen J (1994) *Usability engineering*. Morgan Kaufmann, San Francisco
32. Nielsen J (1997) The changes in Web usability since 1994. *Alertbox* 1 December.
<http://www.useit.com/alertbox/9712a.html> [accessed March 6 2003]
33. Nielsen J (1999) The top ten new mistakes of Webpage design. *Alertbox* 30 May.
<http://www.useit.com/alertbox/990350.html> [accessed March 6 2003]
34. Nordlie R (1999) "User Revelation" – a comparison of initial queries and ensuing question development in online searching and human reference interactions. In: *Proc. 22nd annual international ACM SIGIR conference on research and development in information retrieval*, Berkeley, CA, pp 11–18
35. North C, Shneiderman B (1997) *A taxonomy of multiple window coordinations*. Department of Computer Science Technical Report #CS-TR-3854, University of Maryland, College Park, MD
36. Paynter G, Witten I, Cunningham SJ, Buchanan G (2000) Scalable browsing for large collections: a case study. In: *Proc. 5th ACM conference on digital libraries*, San Antonio, TX, pp 215–218
37. Rauber A, Merkl D (1999) The SOMLib digital library system. In: *Proc. 3rd European conference on research and advanced technology for digital libraries*, Paris, pp 323–342
38. Salampasis M, Tait J, Bloor C (1998) Evaluation of information seeking performance in hypermedia digital libraries. *Interact Comput* 10:269–284
39. Shneiderman B (1984) Response time and display rate in human performance with computers. *ACM Comput Surv* 16(3):265–285
40. Spink A, Bateman J, Jansen BJ (1999) Searching the Web: a survey of excite users. *Internet Res Electron Netw Appl Policy* 9(2):117–128
41. Todd P, Benbasat I (1987) Process tracing methods in decision support systems research: exploring the black box. *MIS Q* 11(4):492–512
42. Theng YL, Jones M, Thimbleby H (1995) Reducing information overload: a comparative study of hypertext systems. *IEE Digest* 95/223: Colloquium on information overload, London, 6/1–6/5
43. Witten I, McNab R (1994) *The New Zealand Digital Library: collections and experience*. *Electron Libr* 15(6):495–504
44. Witten I, Moffat A, Bell T (1994) *Managing gigabytes: compressing and indexing documents and images*. Van Nostrand Reinhold, New York
45. Xu J, Croft W (2000) Improving the effectiveness of information retrieval with local context analysis. *ACM Trans Inf Syst* 18(1):79–112