

A User Evaluation of Hierarchical Phrase Browsing

Katrina D. Edgar¹, David M. Nichols¹, Gordon W. Paynter²,
Kirsten Thomson¹, and Ian H. Witten¹

¹ Department of Computer Science, University of Waikato, Hamilton, New Zealand
{kde2, dmn, kthomson, ihw}@cs.waikato.ac.nz

² University of California, Riverside, CA, USA
paynter@ucr.edu

Abstract. Phrase browsing interfaces based on hierarchies of phrases extracted automatically from document collections offer a useful compromise between automatic full-text searching and manually-created subject indexes. The literature contains descriptions of such systems that many find compelling and persuasive. However, evaluation studies have either been anecdotal, or focused on objective measures of the quality of automatically-extracted index terms, or restricted to questions of computational efficiency and feasibility. This paper reports on an empirical, controlled user study that compares hierarchical phrase browsing with full-text searching over a range of information seeking tasks. Users found the results located via phrase browsing to be relevant and useful but preferred keyword searching for certain types of queries. Users' experiences were marred by interface details, including inconsistencies between the phrase browser and the surrounding digital library interface.

1 Introduction

The central mission of any digital library service is to make information readily available to users. In practice, people generally seek information by directed search, or serendipitous browsing, or some combination of the two. Typical search interfaces, like those found in digital libraries and Internet portals, allow users to enter queries and retrieve sets of relevant documents. Typical browsing interfaces, often presented via the now-ubiquitous Web hypertext browser, help users navigate between resources in an unplanned, undirected manner; new resources are discovered along the way. Between these extremes lies a wide range of interactive interfaces based on both manually-generated and automatically-created information and supporting varying levels of user knowledge and direction.

Effective information-seeking schemes embody mechanisms that prevent users from drowning in information. Most digital libraries let users browse lists of document metadata such as titles and authors (e.g. [19]). As collections grow, however, raw metadata rapidly becomes too voluminous to scan effectively [15]. The standard solution is to provide a topic-oriented hierarchy—such as a library classification scheme—that permits users to drill down from broad groups of items to more manageable subsets. Hierarchical classification schemes have been refined over decades into invaluable search tools, and are universally used in traditional libraries as the ba-

sis for the logical and physical organization of library holdings. Clearly, high-quality classification information should be used wherever it is available to assist users in their browsing activities.

But manual classification is costly. In many digital library or Web-based document collections, accurate subject categories do not exist, and are not feasible to produce. Machine-readable subject thesauri are useful tools for exploring document collections topically, but in many cases digital library documents are not tagged with thesaurus metadata. The problem may ultimately be solved by automated hierarchical classification, an active research topic with some promise for the future (e.g. [4, 10, 11]). Today, however, alternative access mechanisms are sorely needed.

One alternative to manual classification is to automatically create a subject browser that resembles a thesaurus but is based on the occurrence of interesting phrases in documents rather than authoritative subject classifications. We call this strategy *hierarchical phrase browsing* because it allows users to browse through the phrases that occur in a large document collection, examining the terminology that is actually used in the documents, exploring the context in which phrases are used, and retrieving information based on the presence or absence of phrases. Lexical inclusion—whether one phrase occurs as a sub-phrase of another—provides a natural, intuitive, foundation for the hierarchy.

This paper documents the first formal user evaluation of a hierarchical phrase browsing interface based on automatically generated phrases. We focus on a particular browsing tool whose evolution has been described in the literature [13, 15]. Section 2 outlines related work and describes how phrase browsing differs from conventional search and browsing techniques. We then briefly sketch the operation of the target system, *Phind* (for “phrase index”), from a user perspective. The experimental conditions are detailed in Section 4 and the findings presented in Section 5. Finally we discuss how our findings might influence the design of future phrase browsing interfaces for digital libraries.

2 Hierarchical Phrase Browsing

By *hierarchical phrase browsing* we refer to any means of exploring the terminology used in a large document collection, independent of the particular documents in which terms occur, and based on words and phrases extracted automatically from the target documents themselves. Although lexically based, these phrases constitute a plausible, easily-understood, hierarchical, topic-oriented structure for the information in the collection. Several such phrase browsing interfaces are described in the literature, but none have been subjected to rigorous user evaluation. Previous research has concentrated on system description [1, 7, 15, 18], the quality of assigned index terms [14, 17] and evaluation for efficiency [2]; the contribution of this paper is to describe how users actually interact with a phrase browsing interface.

Gutwin *et al.* [7] describe an interface that lets users explore relationships between keyphrases extracted from all documents in a collection. In a small user evaluation, they found evidence that “a phrase-based approach to indexing and presentation offers better support for browsing tasks than traditional query engines” because it provides

“a different level of information that is more appropriate for some kinds of tasks, and is certainly not detrimental in others.” More specifically, they found that the interface performed best when the initial query is general, since a broader range of possible extensions becomes available; that phrase quality was a significant factor; that some legitimate topics are not well-represented as phrases; and that despite the convenience of the interface, some participants lacked confidence in its results. These results have spurred research into keyphrase quality [8, 9] and the development of new phrase extraction methods [5].

Others have evaluated automatically-extracted phrases and phrase hierarchies. Sanderson and Croft [16] describe a method for deriving phrase hierarchies based on phrase co-occurrence, called *subsumption hierarchies*, which are presented to users through standard menu systems, and evaluated by asking human subjects to identify the relationship between pairs of phrases. They conclude that the hierarchies “emulate some of the properties of manually generated subject hierarchies.” In a later evaluation, Lawrie and Croft [12] compare subsumption hierarchies to other hierarchies by estimating the number of selections a user must make to browse from the top of a hierarchy to a relevant phrase, and observe large differences. Paynter *et al.* [14] compare automatically-extracted phrases (generated by Phind, the same system studied in the present paper) with a manually-created subject thesaurus, and discovered enough overlap to consider using the extracted phrase set as a source of suggestions for the thesaurus maintainers. Wacholder *et al.* [17] describe a *dynamic text browser* system for navigating index terms extracted from documents with natural language processing techniques, and evaluate it based on three criteria: index term coherence, coverage of document content and usefulness. None of these evaluations attempts to assess the user’s experiences with a phrase browsing interface.

3 An Interface for Hierarchical Phrase Browsing

The Phind phrase browsing interface is part of the Greenstone digital library software [19]. Greenstone is a complete system for creating, managing, and distributing digital library collections that is freely available from the New Zealand Digital Library Project (<http://nzdl.org>). Phind browsers can be added to any Greenstone collection, and applied to any combination of the documents’ text and metadata fields.

3.1 Using Phind

Fig. 1 shows Phind in use with a Greenstone collection called *FAO on the Internet*, described in more detail in Section 4. The user enters an initial term in the search box at the top. On pressing the ‘Search’ button, the upper panel appears. This shows the phrases at the top level in the hierarchy that contain the search term—in this case, the word *forest*. The list is sorted by phrase frequency; on the right is the number of times a phrase appears, and preceding that is the number of documents in which it appears.

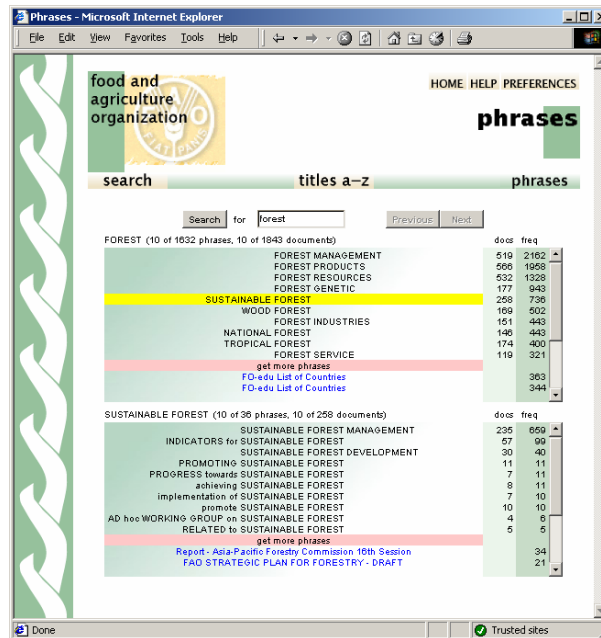


Fig. 1. Browsing for information about *forest* using Phind

The results are initially limited to the first ten phrases because many of the phrase lists are very large. The total number of phrases appears above the list: in this case 10 phrases are displayed of an available 1632 top-level phrases that contain the term *forest*. At the end of the list is an item that reads ‘get more phrases’ (displayed in a distinctive color). Clicking it downloads a further ten phrases, which will be accumulated in the browser window so that the user can scroll through all phrases that have been downloaded so far.

The lower panel in Fig. 1 appears when the user clicks one of the phrases in the upper list. In this case the user has clicked *sustainable forest* (which is why that line is highlighted in the upper panel), causing the lower panel to display phrases that contain the text *sustainable forest*. The text above the lower panel shows that the phrase *sustainable forest* appears in 36 larger phrases, and in 258 documents.

If one continues to descend through the phrase hierarchy, longer and more specific phrases will be found. The page holds only two panels, and when a phrase in the lower panel is clicked the contents of that panel move up to the top panel to make way for the phrase’s expansion in the lower panel. In Fig. 2, for example, the user has expanded *sustainable forest management*, and begun scrolling through its expansions.

The interface not only presents the expansions of the phrase, it lists the documents in which the phrase occurs. Each panel shows a phrase list followed by a document list. The first ten document titles are loaded immediately, and become visible when the list is scrolled. In the lower panel of Fig. 2, the user has scrolled down so that the first six document titles are visible. Document titles are easily distinguished on the screen because they appear in a different color from phrases. On the black-and-white rendition in Fig. 2 they are distinguished by the absence of a “document” count, be-

cause this, by definition, is equal to 1 for the phrase in question (otherwise the document would appear not under the phrase itself but under an expansion of it.) Only the first ten document titles are downloaded, and (as with phrases) users can ‘get more documents’ by clicking on a special entry at the end of the list (which would become visible if the panel were scrolled down a little more).

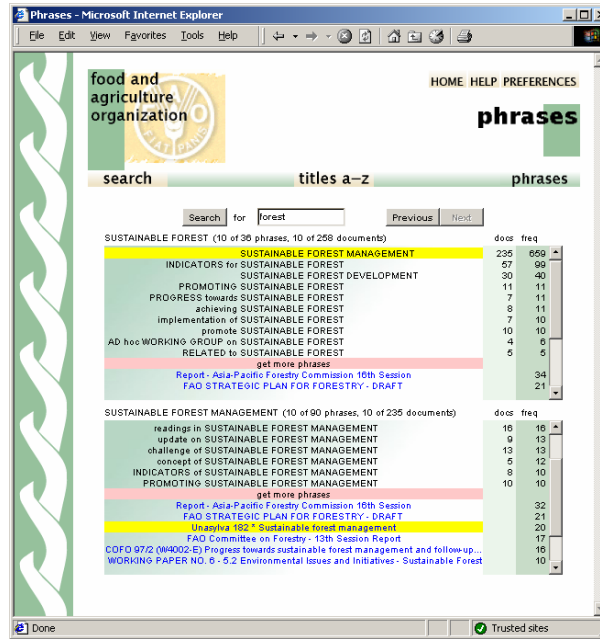


Fig. 2. Expanding on *sustainable forest*

Clicking on a document title opens that document in a new window. In fact, in Fig. 2 the user has clicked on *Unasylya 182 * Sustainable forest management* (that is why its line is highlighted), which displays the document in a new window. As Fig. 2 indicates, that document contains 20 occurrences of the phrase *sustainable forest management*. Greenstone will optionally underline each occurrence of the phrase in the document text.

3.2 Extracting Phind Phrases

Underlying the Phind user interface is a hierarchy of phrases that appear in the document collection. In their evaluation of a hierarchical keyphrase interface, Gutwin *et al.* [7] observed that the user experience is affected by phrase set quality, and by topics that are not well-represented by a browsable phrase. For this reason, the Greenstone version of Phind is based on a completely new phrase extraction scheme that eschews highly selective phrase sets, and instead creates a hierarchy from every repeated phrase in the document collection (with some caveats). As an example, consider the

following paragraph from a document in the *FAO on the Internet (1998)* collection, which is annotated with parentheses to show the phrase structure:

The capacity of the (Desert Locust) for (rapid multiplication) and ((long-distance) migration) ((under [favourable) conditions]) results in highly (variable population) levels over (space and time). Periodically, major population ((upsurges and [plagues) develop]), usually following a sequence of (good rains) which favour reproduction. Such (upsurges and plagues) are interspersed by (periods of (relative inactivity)) called recessions during which ((Desert [Locust) populations]) may remain at very (low levels) for a number of years.

We use the term “phrase hierarchy” to describe the subphrase relationship, though it is not a tree structure; square brackets appear in the above excerpt to indicate non-nested phrases. In Phind, a “phrase” is defined as a sequence of words that occurs more than once in the text. To include *every* such phrase would clutter the interface with trivial phrases, so we add three further conditions to the definition. Phrases must not contain *phrase delimiters*, must neither begin nor end with a *stopword*, and must be *maximal-length*.

The first two restrictions are purely syntactic. If the text were treated as an undifferentiated stream of words, many of the phrases extracted from it would cross syntactic boundaries. To take an extreme example, the last word of one document and the first word of the next are unlikely to form a meaningful two-word phrase. For this reason, we impose the constraint that phrases may not include delimiters. Delimiters are defined as the end of documents, the end of sentences, and any punctuation characters. In practice, we tune the punctuation rule to account for common (and language-dependent) usage: in English, for example, neither the apostrophe in *don't* nor the hyphen in *language-dependent* are interpreted as phrase boundaries. We also mandate that stopwords (like *the*, *of*, and *for*) may not appear at the beginning or end of a phrase to avoid trivial expansions like *the capacity* or *locust for*.

The requirement that phrases are maximal-length sequences is more complex. A phrase is maximal-length if it occurs in more than one context, where by “context” we mean the words that flank the phrase where it appears in the text. Phrases that are not maximal-length—ones that are flanked by the same words wherever they appear—are expanded to encompass the identical context. In the FAO collection, for example, the phrase *forest industries strategy* occurs only in the longer phrase *forestry industries strategy study*, so the latter term is displayed at the top level of the hierarchy in place of the former. On the other hand, the phrase *sustainable forest* occurs in many different contents—ten examples can be seen in the lower pane of Fig. 1.

The three conditions are carefully chosen to provide a phrase structure that covers all the topics appearing in the documents and which makes the phrase/subphrase relationship both simple and obvious. This phrase structure allows an interface to clearly situate a user in the hierarchy, yet it remains simple enough to be generated for large collections with limited computational resources.

4 Study Description

The user evaluation assessed the usability and utility of the Phind phrase browsing interface. The goals of the evaluation were threefold: (i) to determine the value of using Phind to browse a topic as a way of learning what a collection has about that topic; (ii) to determine whether Phind is better than keyword searching in terms of ease of locating specific information held within a collection and ease of navigation; and (iii) to assess the participants' subjective acceptance of Phind's interface.

The study was conducted during February 2003 at the University of Waikato Usability Laboratory. Each session involved only one participant, who performed two tasks during a single session. Because the study was comparative, we used a within-subject design: each participant worked with both the phrase browsing interface and keyword searching. The design was counter-balanced by randomizing the order of the tasks to reduce the effect of transfer learning.

There were twelve participants in the study. All were students (seven graduates) who were nearly evenly split between computing and management disciplines. All but one of the participants typically used computers for more than an hour a day, all were familiar with Internet keyword searching but only nine had previously used a digital library. Three-quarters of the participants were male.

The study used the FAO Collection within Greenstone; it comprises the Web site of the Food and Agriculture Organization (FAO) of the United Nations, in a version that was distributed on CD-ROM in 1998. This is not an ordinary, informally-organized Web site. The mandate of the FAO is to distribute agricultural information internationally, and the information included is controlled, giving it more of the characteristics of a typical digital library collection. With 21,700 Web pages, as well as around 13,700 associated files (image files, PDFs, etc.), it corresponds to a medium-sized collection of approximately 140 million words of text. The Web site (<http://www.fao.org>) has since grown to many times this size, but we use the 1998 version as it was selected by editors at the FAO, and contains no dynamic content.

Each participant was asked to complete seven tasks that involve locating information, understanding content, and recognizing and using elements and functions. Participants were asked to complete a variety of information searching tasks with both the usual Greenstone keyword searching and the Phind interface. Participants were not given any explicit training but were prompted with help during their first task. The study included both exploratory questions, e.g.:

–‘find out more about national forest programmes in different countries’

and specific retrieval tasks, e.g.:

–‘where can golden apple snails be found?’

–‘what was the locust numbers situation during May in Kuwait?’

–‘what does APFSOS stand for?’

The participants used a browser to access the FAO Collection on the New Zealand Digital Library Web server. The Phind interface is presented as a Java applet within Greenstone. Participants' interactions were recorded on video and they were also

asked to complete questionnaires before and after the tasks, and a final comparison questionnaire.

5 Results

The analysis of the session video recordings and the participants' questionnaires yielded an abundance of data; we present the main findings here in terms of task performance, general usability of the Phind interface and usability problems discovered in the interface.

5.1 Task performance

Most participants (10 out of 12) indicated in their questionnaire that they had found the results returned by Phind to be *clear and easy to understand*. Similarly, 10 out of 12 indicated that Phind's results were *relevant and useful to the query*. For both these questions the participants rated Phind to be clearly superior to keyword searching.

Exploratory Questions. The participants were generally positive about performing exploratory tasks in Phind. These tasks gave rise to positive feedback regarding the way the system is set out and subphrases are presented.

Some participants spent in excess of 10 minutes completing these questions. Its exploratory nature meant they were prepared to keep going almost indefinitely, finding out more detailed information than was required. Most participants used this time to explore the interface, and some asked questions regarding the actions that occurred. This period exposed many of the usability issues reported in Section 5.3.

Most users followed the same paths through the top level of the hierarchy on multiple occasions (e.g. *Forest*, then *national forest*, then *national forest programmes*), before following different avenues with more specific phrases, which suggests that many participants were unaware of the function of the navigation buttons.

Specific Retrieval Tasks. The participants had difficulties using phrase browsing to complete the specific retrieval tasks that involved multiple topics, e.g. 'what are the most widely planted pines for timber and pulp production in the southern United States?' and 'What was the locust numbers situation during May in Kuwait?' Of the 12 participants who attempted these two questions with Phind, four gave up, five gave the wrong answer and only three found the correct answer. In stark contrast, these two tasks were successfully completed by 11 of the 12 participants who used keyword searching.

5.2 The Phind Interface

Users readily accepted the idea of phrase browsing; in the summary questionnaire two thirds of participants listed the concept of Phind as the *element or feature that they liked the most*. Comments included: “The way you filter through the results to narrow it down” and “The idea of it – takes the work out of searching for you.”

Three-quarters of the participants commented on Phind’s inability to search on more than one term. All participants tried at least one multi-term search in Phind. *Being able to search on only one term* appeared in the summary questionnaires of three participants as an element or feature that they disliked most. Comments included: “You should be able to put more than one word” and “Confusing when I was searching for two different topics.”

5.3 Phind Usability Issues

Analysis of the participants’ interactions, particularly in response to the exploratory questions, identified a range of usability problems in the Phind user interface.

Two Window Display. The use of two windows was problematic for some participants. Three participants minimized the document window instead of closing it, which meant that when they clicked on a document link, Phind opened the document in the hidden window. This was not apparent to the participants, who received no indication at all that the document had opened. All three of these participants had to be prompted repeatedly that their document could be found in the other window.

Results Panels. Five participants had problems with the relationship between Phind’s two result panels. For example, one thought that the lower panel was always a subsection of the upper one. Three thought that the phrases would remain in the upper panel and the documents would remain in the bottom one. Confusion arose when these participants investigated a second phrase link from the top screen and ended up with the results of their two phrase links being displayed and their list of phrases in the history.

Title Display. The titles of the returned documents are centered in the display space (see Fig. 2 above for an example). However, when the document title exceeds the width available it is truncated at both ends, and users have no means to scroll horizontally. Half the participants commented negatively on this truncated display. Often the displayed portion of the title is insufficient to indicate a document’s topic. Comments included: “It would be nice if you could read at the sides” and “How are you going to know what’s contained in a paper if you can’t see the title?”

At the phrase link level of the hierarchy the same truncation occurs, but with slightly different results. This is because phrase links are aligned on the search term that they contain, rather than on the text as a whole.

User Feedback from Requests. Five participants found it frustrating that there was no feedback after a request to indicate that processing was occurring. Impatient partici-

pants often clicked a link several times while waiting for Phind to load the next results page. To prevent this re-occurring, some participants were eventually told verbally that Phind was indeed processing their request.

Navigation. The navigation provided by Phind was underused. Five of the twelve participants did not use the 'Previous' or 'Next' buttons at all. One did not appear to understand the 'Previous' button and listed "not being able to go back" as one of the *elements or features that they most disliked about Phind* in their summary questionnaire. Of the seven participants who did use the navigation buttons, two participants first had to enquire, "Is there a way to go back?" We also observed that few participants followed the 'get more phrases' and 'get more documents' links, suggesting that their function was poorly understood by some participants.

5.4 Summary

Although two-thirds of the participants liked the concept behind Phind, only a quarter named it as the search method they preferred overall. The two main reasons appear to be difficulties with multi-term topics and queries, and assorted usability issues with the interface.

6 Discussion

Users readily accepted the phrase browsing concept, despite several shortcomings in the Phind interface, and were able to understand the results. In this sense, the interface is a success.

However, three quarters of the users preferred the keyword searching over phrase browsing overall. Despite liking the Phind interface, the participants found many problems. The main functional problem was Phind's inability to perform multi-word queries. Another problem is Phind's unfamiliarity: with a new interface like Phind, a learning period is required, and the limited verbal help we provided is unlikely to match the experience of our participants already had with keyword searching. On the other hand, this unfamiliarity may lend an aspect of novelty value to its appeal.

The study showed that the issue of locating a starting point for a phrase browsing interaction was problematic. In providing a text box for a user to enter a term the interface suggested that any terms could be entered; as in, say, a web search engine. A common result of this is the zero-hit response as the user's terms did not exist as a phrase in any of the documents. However this behaviour contrasts sharply with browsing once the user has entered the phrase hierarchy, where *all* options displayed are actual phrases from the texts.

The study clearly implies that the Phind interface should be refined to distinguish specific interface effects from more general phrase browsing issues. Although the usability issues were individually small, they clearly had a significant impact on the participants' experiences. In this regard the study reinforces existing usability studies of digital library use. In particular, we observed two previously-reported design issues:

“working across boundaries” in the different paradigms of browser-based keyword searching vs. the Java-based Phind interface; and “blind alleys” when Phind users attempted multi-term phrase queries [3]. ‘Boundary’ related problems included inconsistent experiences with the opening of windows leading to lost documents, lack of feedback during query evaluation, unfamiliar navigation tools, and problems understanding the relationship between frames and result sets.

Phind is implemented as a Java applet within the Greenstone digital library, which is otherwise HTML-based. Several issues appear to have arisen because Phind was originally developed as a stand-alone application connecting to a customised server [13] and was only relatively recently ported to the hypertext environment. We hope that by making Phind more consistent with Greenstone’s usual behaviour, and with the users’ usual expectations of Web site design, these problems will be resolved.

However, there is an alternative perspective. It is attractive to consider re-casting Phind as an interface in which it is impossible to express zero-hit queries [6]. If the initial word were selected from a list of the actual vocabulary of the document collection instead of being freely typed, users could drill into the contents of the collection simply by clicking on words and phrases to reveal ever longer ones. This pulls the implementation strategy away from HTML and firmly into a more reactive interface, revealing a tension between trying to conform with existing search interface paradigms and pushing the envelope of reactive interfaces.

There are clearly usability issues with the Phind interface, but the participants’ reactions to phrase browsing are encouraging. In addition to interface design work, a larger longitudinal study would be interesting to examine whether these positive reactions transfer from the usability laboratory to everyday information searching.

7 Conclusion

Although hierarchical phrase browsing systems have existed for some time, there appears to be little research on users’ experiences with this style of interaction. This study compared a stable, familiar search interface with a novel hierarchical phrase browsing system. The upshot was that although the participants preferred the search interface, they could see the potential advantages of phrase browsing.

Through its universal adoption in Internet search engines, the keyword search paradigm dominates information retrieval today. Phrase browsing interfaces promise something more—an interaction with greater structure that offers improved support for users’ exploration of document collections.

References

1. Anick, P., Tipirneni, S.: The Paraphrase Search Assistant: Terminological Feedback for Iterative Information Seeking. In Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. Berkeley, CA, USA (1999) 153-159

2. Bahle, D., Williams, H.E., Zobel, J.: Efficient Phrase Querying with an Auxiliary Index. In Proceedings of the ACM-SIGIR Conference on Research and Development in Information Retrieval. Tampere, Finland (2002) 215-221
3. Blandford, A., Stelmaszewska, H., Bryan-Kinns, N.: Use of Multiple Digital Libraries: A Case Study. In Proceedings of the First ACM/IEEE-CS Joint Conference on Digital Libraries. Roanoke, VA, USA (2001) 179-188
4. Dolin, R.: Pharos: A Scalable, Distributed Architecture for Locating Heterogeneous Information Sources. PhD Thesis, University of California, Santa Barbara, Santa Barbara, CA, USA (1998)
5. Frank E., Paynter G.W., Witten I.H., Gutwin C., Nevill-Manning C.G.: Domain-specific Keyphrase Extraction. In Proceedings of the International Joint Conference on Artificial Intelligence. Stockholm, Sweden (1999) 668-673
6. Greene, S., Tanin, E., Plaisant, C., Shneiderman, B., Olsen, L., Major, G., Johns, S.: The End of Zero-Hit Queries: Query Previews for NASA's Global Change Master Directory. International Journal of Digital Libraries 2(2-3) (1999) 79-90
7. Gutwin C., Paynter G., Witten I.H., Nevill-Manning C., Frank, E.: Improving Browsing in Digital Libraries with Keyphrase Indexes, Journal of. Decision Support Systems 27(1-2) (1999) 81-104
8. Jones S., Paynter G. W.: Automatic Extraction of Document Keyphrases for Use in Digital Libraries: Evaluation and Applications. Journal of the American Society for Information Science and Technology 53(2) (2002) 653-677
9. Jones S., Paynter G. W.: An Evaluation of Document Keyphrase Sets. Journal of Digital Information 4 (1) (2003)
10. Koller, D., Sahami, M.: Hierarchically Classifying Documents Using Very Few Words. In Proceedings of ICML-97: the 14th International Conference on Machine Learning. Nashville, TN, USA (1997) 170-178
11. Larson, R. R.: Experiments in Automatic Library of Congress Classification. Journal of the American Society for Information Science 43(2) (1992) 130-148
12. Lawrie, D., Croft, W. B.: Discovering and Comparing Hierarchies. In Proceedings of Recherche d'Informations Assistee par Ordinateur (RIA/O) Conference. Paris, France (2000) 314-330
13. Nevill-Manning, C.G., Witten, I.H., Paynter, G.W.: Browsing in Digital Libraries: A Phrase-based Approach. In Proceedings of the Second ACM International Conference on Digital Libraries. Philadelphia, PA, USA (1997) 230-236
14. Paynter, G.W., Witten, I.H., Cunningham, S.J.: Evaluating Extracted Phrases and Extending Thesauri. In Proceedings of the Third International Conference on Asian Digital Libraries. Seoul, Korea (2000) 131-138
15. Paynter, G.W., Witten, I.H., Cunningham, S.J., Buchanan, G.: Scalable Browsing for Large Collections: A Case Study. In Proceedings of the Fifth ACM Conference on Digital Libraries, San Antonio, TX, USA (2000) 215-223
16. Sanderson, M. Croft, W. B.: Deriving Concept Hierarchies from Text. In Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Berkeley, CA, USA (1999) 206-213.
17. Wacholder, N., Evans, D.K., Klavans, J.L.: Automatic Identification and Organization of Index Terms for Interactive Browsing. In Proceedings of the First ACM/IEEE-CS Joint Conference on Digital Libraries. Roanoke, VA, USA (2001) 126-134
18. Wacholder, N., Nevill-Manning, C.: Workshop 2: the Technology of Browsing Applications. In Proceedings of the First ACM/IEEE-CS Joint Conference on Digital Libraries. Roanoke, VA, USA (2001) 483
19. Witten, I.H., Bainbridge, D.: How to Build a Digital Library. Morgan Kaufmann, San Francisco, CA, USA (2003)

Edgar, K.D., Nichols, D.M., Paynter, G.W., Thomson, K. and Witten, I.H. (2003) A user evaluation of hierarchical phrase browsing, *Proceedings of the 7th European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2003)*. LNCS 2769. Berlin/Heidelberg: Springer. 313-324.

http://dx.doi.org/10.1007/978-3-540-45175-4_29

This is an author-created version. The final publication is available at www.springerlink.com