

Searching Emergent Vocabularies: *Exploring Methods to Reduce Cognitive Load during Web Navigation and Resource Contribution*

Samuel Joseph, Joyce Yukawa, Daniel Suthers and Violet Harada
*Department of Information and Computer Science
University of Hawai'i
{srjoseph,yukawa,suthers,vharada}@hawaii.edu*

Abstract

This paper presents a differential usage study of a web-based resource database that provides both search and associative browsing functionality. The associative browsing is based on emergent meta-data: meta-data that is derived from the terms that users associate with resources they have contributed to the system. We argue that this approach provides a low cognitive load information seeking mechanism, and can also reduce the effort required by the user to enter meta-data when contributing resources. In this paper we concentrate on a three-month study of student librarians using the system, with analysis of their activities and other data collected by questionnaire. The results suggest that associative browsing was at least as popular as search, and that providing perspectives on emerging meta-data during the contribution process may have helped the community self-organize a vocabulary.

1. Introduction

Computer networks provide the user with access to an enormous breadth of resources, including huge digital libraries and archives. For the individual user, filtering the information available from these sources is a daunting task. Many portals arrange their contents in a hierarchical fashion, which works well for a limited volume of material, but starts to become a barrier as quantities increase [23]. The problem is that the structure of the hierarchy classifies individual items in unique locations, defined by a path from the root of the hierarchy to a particular leaf. As the quantity of material increases, the path to a resource becomes longer and longer.

Digital systems allow the same items to be placed in more than one location within a hierarchy, but the users' ability to find something also depends on a match between their own vocabulary/ontology and

that of the hierarchy designer regarding terms that represent particular concepts and define broader supersets and narrower subsets for each concept. A careful study of the vocabulary/ontology of a particular community can support the creation of a hierarchy that meets the needs of that community at a given point in time [15], but must be continually maintained to meet evolving community needs.

An alternative to hierarchical navigation is to provide search capabilities over a set of resources. A full-text search is one approach, but there is no guarantee that the contents of a document will include the query terms employed by the user [2]. Search over assigned meta-data fields such as keywords can resolve this problem, but this fails to the extent that the keyword choices of the user differ from that of the resource contributor [13], or that descriptive needs change over time [11].

So how can we provide efficient navigation support over large digital libraries that are serving an evolving community? Recently, social bookmarking tools [4] have popularized some alternative approaches. These approaches, such as “associative browsing” and “emergent community meta-data”, have shown promise in addressing the problem of adapting to an evolving community. In “associative browsing”, a user is presented with a set of topics associated with an information query they have just made. It is different from pure hypertext navigation in that the topics presented are generated dynamically from the associations present in some knowledge base. “Emergent community meta-data” may serve to underpin “associative browsing”, but the two may also be considered separately. By itself, “emergent community meta-data” is meta-data generated in a bottom-up fashion through the many isolated actions of individuals within a community. This is in contrast to meta-data that has been generated in a top-down fashion according to some organized schema.

To the extent that the community has some degree of consistency in their meta-data markup activities, one can expect to see emerging trends, such as popular terms and popular sub-sets of terms based on associations with particular sets of resources. Conversely, in the absence of any central coordination, the community vocabulary may diverge or splinter into sub-groups. This is not intrinsically a problem in as much as the meta-data associated with each sub-community can serve the needs of that sub-community, although it might be challenging to present views that distinguish the knowledge space of each group. A more serious problem is that of non-uniform meta-data across the entire collection. This might include spelling variations of the same word or numerous synonyms for a concept. This kind of inconsistency can lead to a resource collection that is populated with a confusing array of overlapping terms and is thus difficult to search effectively. These problems must be overcome if emergent meta-data is to effectively meet information seekers' needs. Potential solutions include recommending popular community keywords during resource contribution, and seeding a community with resources marked up according to a controlled vocabulary.

Assuming that the problem of divergence can be solved, many types of knowledge systems could benefit from associatively browsable interfaces and emergent community processes. The main benefit would be to keep knowledge systems synchronized with the needs and preferences of their users. This could apply independently of whether they are digital libraries or archives where the resources are contributed by a select set of specialist users; social bookmarking systems where any community member can contribute a resource; or digital libraries where anyone can submit a resource but the resources themselves will be subject to a review process before addition to the database.

This paper describes the results of a three-month study of graduate students enrolled in a school library practicum course using the Hawai'i Networked Learning Communities (HNLC) web-based resource database. The web access to the database was augmented to provide both associative browsability of community meta-data and meta-data suggestion support during resource contribution. The following subsection considers the concept of associative browsability in more detail, while section 2 reviews some related work. Section 3 will provide background on the HNLC project and its resource database, and in section 4 we will describe the design of the study and the experimental hypotheses. Subsequent sections will present descriptions and analyses of the results.

1.1 Associative Browsability

Browsability itself (e.g., enabling the user to examine a resource database through lists of collections, topics, or organizations) is not uncommon in digital library interfaces [18]. Associative browsability goes further in that rather than descending a hierarchy of one type or another, the user is shown additional relations based on shared meta-data associations. For example, a user clicking on the topic "electromagnetism" would see not only the resources associated with this topic, but also the other topics that are related to the set of resources related to "electromagnetism."

Of fundamental importance is that documents or other resources can have multiple associations, so that when browsing the documents with a particular association the user can see further associations of the resulting documents that support both narrowing and broadening of the knowledge seeking process. For example the most frequent terms that are associated with electromagnetism might include: "Physics", "Magnetism", "History", "Maxwell Equations", "Encyclopedia", "Charges", "Lecture", "Software Design", "Optics" and "Faraday". If the user is keen to narrow the search focus they can add one of these terms to their query, or if they are interested in broadening or changing the focus they can click through to see all resources related to physics or optics. Naturally there is a burden on the interface designer to make it clear to the user how terms can be clicked through in such a fashion as to narrow or broaden the search. It would also make sense to provide additional information to support the user's choices, such as the number of resources that will be returned by a branch out query to "Physics", or a more specific one to "Electromagnetism" and "Faraday".

Arguably this type of associative browsing is not dissimilar to the process afforded by many classification schemas, such as the Library of Congress Classification System and numerous specialized subject thesauri. Typically, after querying a database, the user is presented with other terms associated with each resource retrieved, providing the opportunity to browse using those other terms. In some systems, such as Cambridge Scientific Abstracts, the user may choose to browse keyword lists (alphabetical or hierarchical lists or a rotated index of related terms) separately from querying the databases to retrieve resource documents. However, these systems rarely present the broadening/narrowing options dynamically in response to the current query.

In addition, combining associative browsability with community meta-data adds a dynamic

community perspective, such that rather than browsing an established hierarchy, the user is now browsing a knowledge space that has emerged from the aggregate activities of multiple community members. Contributing members of the community now have a chance to influence this knowledge space, to have their own needs responded to directly. Many different perspectives on the emerging community meta-data are possible, such as frequent, popular or recent terms used for searching or annotating resources. It is an open question as to which perspective or collection of perspectives will best serve the needs of the users and help them contribute successfully to the community.

2. Related Work

The combination of associative browsing with emergent community meta-data is relatively recent, making it hard to place them in the context of the scientific literature, but they can be thought of as a type of “social navigation” [8]. Social navigation in this context is a process whereby the navigation of an individual is supported by the aggregated actions of other community members, such as navigating to those resources most popular within the community. One example of this is the related books feature of Amazon.com, “customers who bought this book also bought ...”. Although this particular service is a type of collaborative filtering [22], emergent community meta-data distinguishes itself from collaborative filtering by its larger scope. Collaborative filtering is usually restricted to shared purchases or scalar ratings of a product's quality. Emergent community meta-data considers a much more complex markup of an individual resource.

Systems that support this emergent community metadata can lead to the creation of a browsable “folksonomy” [17], a dynamic taxonomy that represents the categories that individual users are employing to organize their own information spaces. Mathes [17] describes how a folksonomy is more similar to a categorization than a classification. A categorization emphasizes synthesis of similarity while classification emphasizes systematic arrangement of materials [12]. In a folksonomy each document can have many terms associated within a flat namespace, in contrast to classification schemes that tend to be hierarchical and provide a single unique classification for each item.

Excellent examples of folksonomies can be seen in social bookmarking sites such as Delicious and Flickr that provide views of emergent meta-data as users tag bookmarks and photos (respectively) with arbitrary tags or keywords of their own devising. Hammond et

al. [10] provide an overview of these and other systems. In the scientific literature the closest approaches are the document popularity navigation system employed in KnowledgeSea [3] and the web query mining approach of Davison et al. [6].

General-purpose digital libraries appear to focus on search interfaces and automating parts of the resource contribution process, e.g. INFOMINE [16], a tiered system where entries are graded according to whether they have been edited/created by an artificial expert classifier or a human expert and additionally by whether the resource experiences high usage and linkage. There appears to be no dynamically created perspective on the wide range of available meta-data, and the popularity of the different resources is not displayed as a browsable perspective.

There are efforts underway to enhance existing digital libraries through graphical browsing of results [14], browsing by virtual shelves [21], automatic thesaurus generation techniques to create graphs of domain specific concepts [5], and allowing a searcher to pose queries to one viewpoint and then change to another viewpoint while retaining a sense of context [9]. As yet we are unaware of any attempt to add a community meta-data dimension to a digital library.

3. HNLC

Hawai'i Networked Learning Communities (HNLC) is a cooperative effort by the Hawai'i Department of Education, the University of Hawai'i Information and Computer Sciences Department, many community partners, and the teachers and students at rural public schools throughout the Hawaiian Islands. The web site (hnlc.org) provides online access to a professional community of educators, including news and discussion forums, a resource database (RDB) and HNLC teacher-developed unit plans.

Teachers develop and share inquiry-based lessons in math, science and technology that are rooted in the local environment and culture. The HNLC RDB provides web links and information on curricula and curriculum support materials, educational software programs, environmental education projects, and places throughout the islands that can support project-based learning. HNLC participating teachers can also discuss each locally contributed resource with its author and others - how to use it, how to adapt it, and what experiences they have had. These resources can be browsed via keyword or accessed using a search form. Any HNLC member can submit links and online resources to help expand the coverage of resource listings.

In addition teachers have privileged access to many parts of the site and can use an HNLC bookmarking system to store references to other members, resources, unit plans and so forth. A member's bookmarked items then appear in the member's personal space within the site.

4. Study Design

As part of our efforts to improve the service provided by the HNLC RDB we implemented support for community meta-data and conducted a three-month study involving graduate students enrolled in a school library practicum course. The purposes of this study were to assess the usability of the new features and to investigate some of the consequences of their introduction.

Our central research questions were: (1) Does associative browsing of emergent community meta-data support information seekers' needs? (2) How can the problem of vocabulary divergence be addressed? In order to answer these questions we observed patterns of use of the RDB as participants both bookmarked items and posted additional resources. The intention was to analyze the patterns of behaviour and elicit participant responses regarding the usability of the system.

The participants took part in a training session about how to post resources, how to navigate through the RDB, and how to bookmark resources. Ten participants took part in the study that ran from February to April 2005. The participants were all graduate students who were asked to post one new resource and bookmark two existing resources each month of the study.

The experimental hypotheses of the study were as follows:

H1: Users will prefer to seek resources by associative browsing as opposed to employing a free text search form.

H2: The vocabulary of a community will diverge less if popular community keywords are recommended during resource contribution.

The first hypothesis is based on studies showing that different interfaces give rise to different levels of cognitive load. Cognitive load refers to the total amount of mental activity imposed on working memory at an instance in time [25]. Research showing users' lack of willingness to modify queries or to provide relevance feedback in web searches [1] has been taken to be indicative of a high state of cognitive load. It is suggested that the user is trying to maintain his or her search criteria in working memory while

considering the set of search results, which itself makes demands on working memory. To explicitly show the difference, Dennis et al. [7] used a secondary digit-monitoring task to show that the cognitive load was lower when using query refinement mechanisms than when perusing document summaries. A query refinement mechanism that suggests alternate search terms makes fewer demands on working memory than document summaries, due to its lower level of complexity.

We hypothesize that the cognitive load of formulating a free text search query is greater than that of employing a browsable interface. In order for a user to select their own query term for entry into a search box she must perform a cognitively more expensive memory recall operation, while browsing even a long list of suggested terms requires less costly recognition operations [20]. Also, we predict that users will prefer to avoid using a keyboard to enter text into forms and will find it easier to navigate and organize meta-data through simple mouse-clicks on suggested terms.

Nielsen's [19] usability studies suggest that more than half of all users are search-dominant, while a fifth of the users are link-dominant, and the rest exhibit mixed behavior. Nielsen describes a typical user comment as: *"I don't want to have to navigate this site the way they want me to. I just want to find the thing I'm looking for."* Superficially this might seem to contradict our first hypothesis. However, Nielsen studied site navigation links (often organized hierarchically), while we are discussing community meta-data links that are like suggested search terms. Arguably browsing emergent community meta-data is more like a type of search: a set of ranked results is returned just as if the user had typed in a search term.

The design of the HNLC RDB user interface supports assessing the validity of H1 since it provides both free text search and associatively browsable keywords on the left and right hand sides of the screen respectively (see Figure 1). All participants received full instruction in using both types of resource seeking mechanism. During the three-month study, log data of user activity was stored and could be analyzed for indications of a preference for one type of resource seeking mechanism over another.

The second hypothesis is at least partly derived from the same studies that suggest the first hypothesis. To the extent that formulating a free text search query places a higher cognitive load on a user than clicking on a suggested search term, it follows that typing free-text keywords as part of annotating a contributed resource is more cognitively demanding than selecting suggested keywords from a list, although it is unlikely that all of the desired keywords for a

particular resource could be displayed in short enough lists for practical scanning of the available keywords. A compromise might be keyword completion whereby the user types the first letters of a keyword and is offered possibilities for completed keywords, although the user may miss other terms starting with a different letter. Each mechanism has advantages and disadvantages, but it seems likely that they would reduce the cognitive load of adding meta-data to a resource, thus encouraging the user to make selections from the suggested keywords. Thus we have a potential mechanism for increasing the likelihood of certain keywords being used. For example the suggested keywords could come from the set of keywords previously used by the user contributing the resource, or they could come from the keywords previously used by other members of the community. It seems likely that users who are offered suggestions from within their own set of existing keywords are less likely to enter keywords with spelling variations and are less likely to choose synonyms to represent

the same concept. The same argument follows for the community at large, although it remains to be seen whether an interface can be devised that would combat divergence over a large community.

However, compared to H1, assessing the validity of H2 was more difficult in this type of study. The approach taken cannot do more than suggest the possible validity of H2 and point towards future studies that might more concretely determine the validity of H2. We took the approach of providing a resource contribution interface that recommended possible keywords that had previously been employed both by the user and by the community as a whole (see Figure 2). We then observed the divergence or convergence of the community vocabulary over the course of the experiment.

As a result of the current design any change in vocabulary convergence could be attributed to other factors such as the natural development of the community. As a consequence the real focus of this paper must be H1.

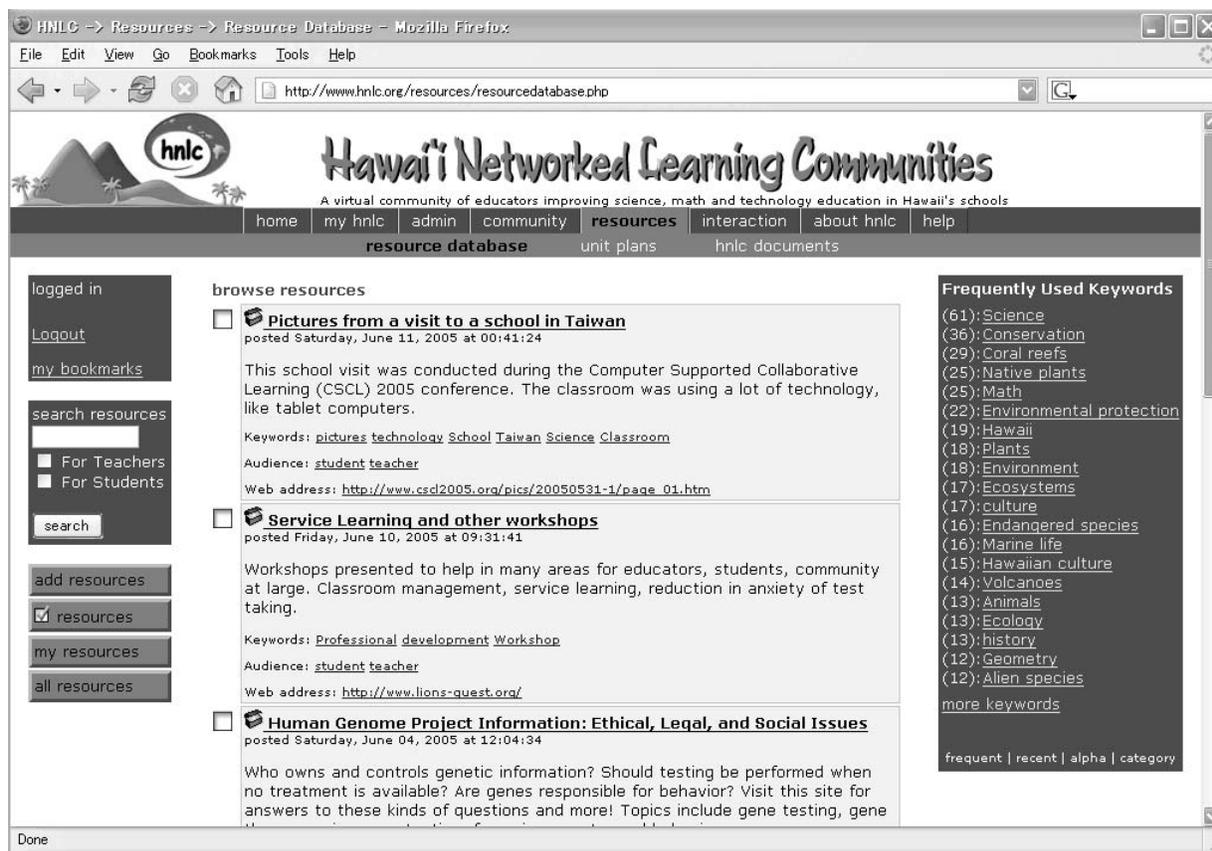


Figure 1. The main HNLC RDB Search Interface; typing free text into the search box on the left or clicking one of the keywords on the right replaces the default reverse-chronologically ordered view with resources ordered by degree of query match. Clicking on the small box to the left of each resource allows the user to bookmark that resource for subsequent retrieval through the “my hnlc” personal pages.

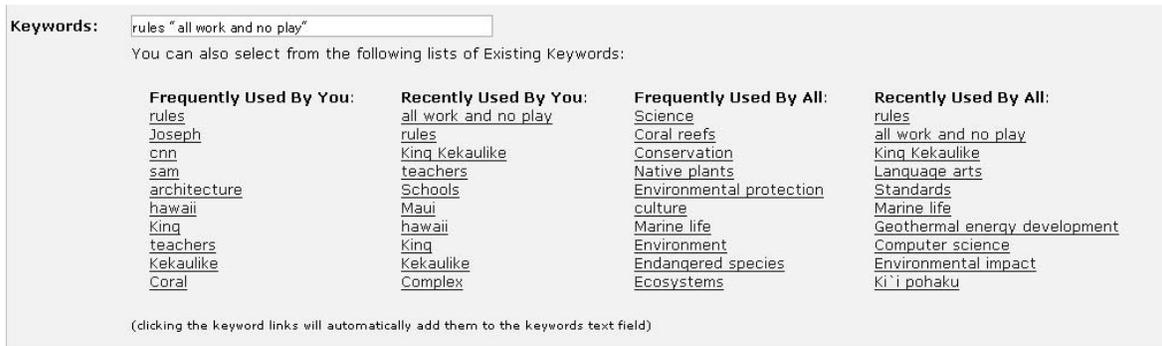


Figure 2. A portion of the HNLC RDB Resource Contribution Interface; clicking on a keyword link adds that keyword to the free text keyword text field. Keyword lists are the top 10 most recently or frequently used for both the user in question, and the community as a whole.

5. Results

All of the 10 participants met the baseline activity requirements and many exceeded them, so that by the end of the study, the 10 participants had contributed 44 resources (30 expected baseline) and 110 bookmarks (60 expected baseline). This had taken place over a total of 2615 page views within the site, and 284 departures from the site to outside resources provided by the RDB.

5.1 Results bearing on Hypothesis 1

After completing the three-month study all participants completed an online questionnaire consisting of 12 questions. The first question was intended to probe the participants preferred information seeking preference, and a summary of the responses is shown in Figure 4, indicating a fairly even split between the two alternatives provided.

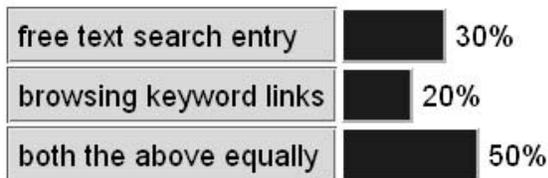


Figure 4. Question 1 results: What is your preferred way of searching?

Analysis of the web site log data backs up this even split, as shown in Figure 5, where 6 of the participants browsed more frequently than they searched, and overall the total number of searches and keyword based browses was almost identical.

While these results do not support H1 directly, it is intriguing to compare these results with those of Nielsen [19] who suggests that 50% of users are searchers, and only 20% link followers. Our study suggests that associative browsing is as popular as simple link navigation, and indeed individual users have widely differing tendencies as implied by several information retrieval models. It is also interesting to consider the possible additional benefits of associative browsing with query refinement. For example, click-through on related meta-data occurred slightly more frequently than query refinement of text based search: 17 (out of 190) query refinements following a meta-browse operation compared to 14 (out of 191) query refinements following search operations.

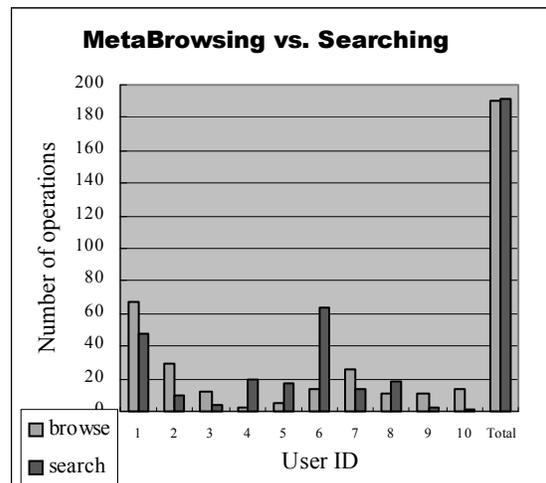


Figure 5. Analysis of search log data

The generally small number of refinements matches larger studies showing that users do not generally refine queries [24]. The difference between the numbers of refinements following search and

browse operations is likely due to random variation. Even if it was not, and subsequent studies show a significant increase in the number of query refinements following browse operations, the result could be interpreted in two ways - (1) browsing operations were less likely to meet information needs thus requiring more query refinements, or (2) query refinement was easier as a result of the associative browsing support.

Of the 284 departures to sites from the RDB, at least 64 of them followed a keyword search, but it would be incorrect to conclude that the remaining 220 departures do follow a free text search. Unfortunately, an oversight in the RDB design made it difficult to effectively categorize the majority of the searches. If the user conducted a search and then clicked on one of the five results returned in the first page, the log entry would include both the departure from the site and the type of search that had originally been conducted. However if the user navigated through to results further down the list, such as results 6-10 or results 11-15 and then clicked through to another site, the log entry would not contain any reference to the type of search that had originally generated this query set. In principle the information could still be extracted from the logs by a more detailed, but extremely time-consuming analysis. Hopefully we may be able to present those results in a subsequent paper that includes a more detailed step-by-step analysis of the users' activity patterns. In addition the oversight has now been corrected and subsequent studies will give a clearer indication as to whether free text or keyword-based searches lead to more "click through" to sites provided by the RDB.

The next two questions in the questionnaire were intended to query the motivations behind using a particular information seeking style. Question 2 was "Under what conditions do you tend to use keywords provided on the site?" A number of respondents described how they used the browsable keyword functionality to aid them during the resource submission process, such as to see which resources were being added under particular topics, and to see what keywords were being used to annotate those particular resources. It was also suggested that the keywords were helpful both for scanning topics, and to see what was available within a particular topic. One respondent noted that the keyword links associated with the display of each individual resource was also a helpful navigation aid, i.e., the ability to click through on the keywords associated with a particular resource and see all the resources associated with that keyword. Given that the associative browsing panel only displayed the top 20 most frequently used keywords, showing the

keywords associated with each resource allowed another route for users to discover different keywords in use within the RDB. Other respondents discussed using keywords to determine how others had marked up particular sites. Although this question was intended to elicit responses related to searching behaviour using keywords, some respondents seem to have interpreted the question more broadly suggesting the need for refinement of this particular question in future studies to distinguish the multiple ways in which keywords can be used within the site: associative browsing, exploring resource meta-data and resource contribution

Question 3 was "Under what conditions do you tend to use your own search terms?" Several respondents noted that most of the time, or at least initially, they started a search with their own search terms. Several others indicated that they used their own search terms when they had a specific topic in mind, but also when they were just browsing and had no specific objective or deadline to meet. Other respondents described how they would fall back on free text search when there was no suggestion (i.e., keywords) matching their particular information need. The rest of the questionnaire questions and results can be found in the appendix.

5.2 Results bearing on Hypothesis 2

Other results came from the log data indicating the use of keywords during resource contribution. Over the 10 participants, 237 keywords were added as meta-data to their 44 resources, giving us an average of just over 5 keywords per resource. In the case where all the users were assigning keywords completely independently we might expect some degree of overlap between their assignments, but let us first consider the extreme case where none of the keywords selected by the users matched those of any other of the users in the community. In this case we would expect the number of unique keywords in the entire community to exactly match the sum of the unique keywords for each individual member, as shown in the following equation, where k = number of unique keywords.

$$\text{community } k = \sum_{\text{members}} \text{member } k$$

So the extent to which the sum of unique keywords for each individual member is higher than the sum of unique keywords in use by the community is an indication of increased shared use of keywords amongst the community members. Thus in a community where the vocabulary is diverging we

would expect the ratio of total unique keywords over the sum of individual unique keywords to be gradually approaching one. Conversely, in a community where the vocabulary is converging this ratio would decrease, and in the extreme case where the vocabularies across all members was identical it would equal one over the number of community members.

In order to see if there was any change in vocabulary convergence we compared the keywords added in the first 6 weeks of the study with those in the second 6 weeks. The log data indicated that 103 keywords were added during the first half of the study, of which 74 were unique, while in the second half of the study 134 were added, of which 98 were unique. The ratio of unique to total keywords does not seem to change substantially from the first half to the second half of the study (0.72 vs. 0.73). However if we sum the unique keywords used by each individual user we discover that there were 91 in the first half and 120 in the second half. Here we see the ratio of total unique keywords over the sum of individual unique keywords to be decreasing (0.88 vs. 0.82), which suggests the vocabulary of this community is converging. This analysis considers the keywords in the two halves of the study completely independently and so this result is not a consequence of there being more words to overlap with in the second half of the study.

Unfortunately the nature of this study is such that we have no comparison, so there is no way to tell if this convergence is a result of the community based keyword suggestion provided during resource contribution or a natural process for a group of individuals with a similar background working on a similar project. However to the extent that the approach itself is sound it appears that we have a simple metric for measuring the divergence or convergence of a community's vocabulary, which we hope to employ in future studies. Thus for the moment we cannot confirm or refute H2.

6. Discussion

Associative browsability such as that provided by the HNLC RDB is attractive because human thought itself is partly associative: as a new concept is considered it can be related to myriad others in a web of associations that support the evolution of new perspectives. However there still exist many questions over the use of such interface techniques.

As noted before Nielsen [19] suggests that 50% of web-users are searchers, and only 20% are link followers. Nielsen himself has done much to emphasize the importance of providing good search

support for web-based portals. That said, Nielsen studied site navigation links, while we have been investigating community meta-data links that are more like suggested search terms. The same argument applies to hierarchical subject headings, which can also be thought of as similar to suggested search terms. Both differ from navigational links that tend to shift the user from one context to another; navigation links lead to a particular portion of a site, rather than a set of items from throughout the site that meet particular information needs.

Browsing subject headings or emergent community meta-data is arguably more similar to searching: a set of search results are returned just as if the user had typed in a search term, except that users did not have to go to the effort of formulating an explicit query, or type it in. Importantly, users get an overview of the contents in the site, and can be guaranteed a non-zero set of search results, something that free text search cannot guarantee.

While it remains to be established whether users perceive associative browsing as closer to search or navigation, our results indicate that for a group of relatively experienced users of search interfaces, community-based meta-browsing was used to a similar extent as free-text search, but not significantly more so, and so we cannot accept H1 at present. However we intend to investigate whether increased familiarity with this type of interface design affects degree of use of community-based meta-data browsing, and whether we will see similar results for less-experienced users.

The current study did not effectively assess H2 although it did demonstrate a metric that could be used for that purpose. A proper assessment of H2 would involve assigning different types of interface to the different subsets of the community, something not naturally supported by a community site. However we intend to employ such a technique in subsequent studies.

Another point to consider is that the users for this study were not the typical users of HNLC (teachers in math, science and technology). Since the graduate students do not directly belong to the community, their behaviour may not be representative. The graduate students were an easy group to track over a particular period; now that we have refined our analytical approach, future studies will examine the resource related activities of the teacher community.

Further experiments of longer duration involving more users are clearly indicated, both to reaffirm our results relating to information seeking, and also to properly assess H2 using our vocabulary divergence/convergence metric.

8. Appendix: Questionnaire Results

Questions 4, 5 and 6 of the questionnaire were designed to elicit feedback about more general aspects of the RDB search interface. In the future we intend to ask similar questions of less experienced users in order to determine user preference for such features as default ANDing, or ORing of search terms (Figure 6). In the current system the queries are default AND, and there is no advanced search form to offer a choice of OR. The question will also be clarified to elicit user preference for AND or OR as opposed to what the users think the system is providing.

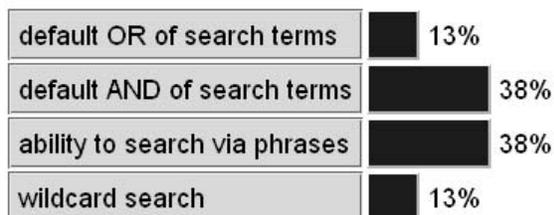


Figure 6. Question 4 results: Which search features do you use most? (Not disjunctive)

The answers to question 5 (Figure 7: “Which aids for browsing would be useful to you?”) indicated that the participants seemed strongly in favor of receiving lists of keywords related to their search terms, a feature a number of online search engines have started providing. The answers to question 6 (Figure 8: “Which way of displaying search results do you prefer?”) confirmed our expectation that users prefer, briefer descriptions of the resources themselves.

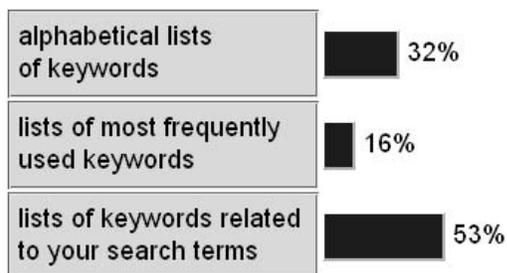


Figure 7. Question 5 results: Which aids for browsing would be useful to you? (Not disjunctive)

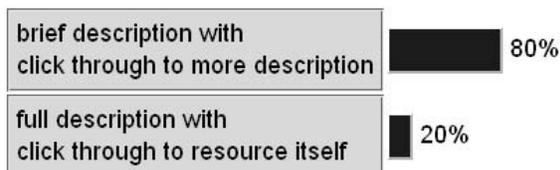


Figure 8. Question 6 results: Which way of displaying search results do you prefer?

The answers to Questions 7 (Figure 9: “When adding resources, what features are helpful?”) and 8 (“What information should be included in a brief description of the resource?”) were somewhat equivocal with no clear preferences for the types of keywords provided during resource contribution, and the types of meta-data that should be included with a brief description of a resource.

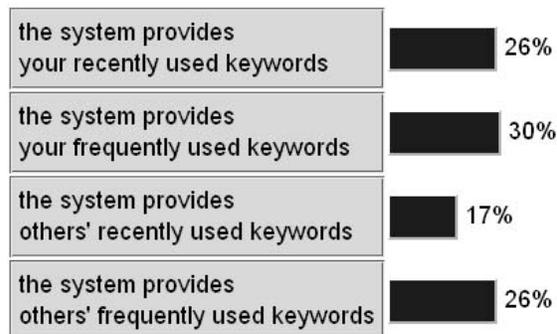


Figure 9. Question 7 results: When adding resources, what features are helpful? (Not disjunctive)

9. References

- [1] Back, J. & Oppenheim, C., “A Model of Cognitive Load for IR: Implications for User Relevance Feedback Interaction”. *Information Research* 6(2), 2001.
- [2] Barnum, C., Henderson, E., Hood, A., & Jordan, R. “Index Versus Full-text Search: A Usability Study of User Preference and Performance”. *Technical Communication Online* 51(2):185-206, 2004.
- [3] Brusilovsky, P., Chavan, G. & Farzan, R., “Social Adaptive Navigation Support for Open Corpus Electronic Textbooks”. In W. Nejdl, & P. De Bra, (Eds.) *Proc. of the 3rd Int. Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, (pp. 24-33). Eindhoven, The Netherlands: Springer LNCS 2004.
- [4] Cayzer, S. “Semantic Blogging and Decentralized Knowledge Management”. *Communications of the ACM* 47(12): 47-52, 2004.

- [5] Chen, H.C., Schatz, B., Ng, T., Martinez, J., Kirchhoff, A., & Lin, C.T. "A Parallel Computing Approach to Creating Engineering Concept Spaces for Semantic Retrieval: The Illinois Digital Library Initiative Project". *IEEE Transactions on Pattern Analysis and Machine* 18(8): 771-782, 1996.
- [6] Davison, B.D., Deschenes, D.G., & Lewanda, D. B. "Finding Relevant Website Queries". In *Proc. of 12th Int. World Wide Web Conference*, (pp308-309), Budapest, Hungary, May 20-24, 2003.
- [7] Dennis, S., McArthur, R., & Bruza P. "Searching the World Wide Web Made Easy? The Cognitive Load Imposed by Query Refinement Mechanisms". In *Proc. 23rd ACM SIGIR Conf. R&D in IR* (pp 280-287), 2000.
- [8] Dieberger, A., Dourish, P., Hook, K., Resnick, P., & Wexelblat, A. "Social Navigation: Techniques for Building More Usable Systems". *Interactions* 7(6): 36-45, 2000.
- [9] French, J.C., Chapin, A.C., & Martin, W.N. "Multiple Viewpoints as an Approach to Digital Library Interfaces". *J. American Society for Information Science & Tech.* 55(10): 911-922, 2004.
- [10] Hammond, T., Hannay, T., Lund, B., & Scott, J. "Social Bookmarking Tools (I) A General Review". *D-Lib Magazine* 11(4): 2005.
- [11] Ives, Z. G., Khandelwal, N., Kapur, A., & Cakir M. "ORCHESTRA: Rapid, Collaborative Sharing of Dynamic Data". *CIDR 2005* (pp 107-118), 2005.
- [12] Jacob, E. K. "Classification and Categorization: A Difference that Makes a Difference." *Library Trends*. Winter, 2004.
- [13] Kitajima, M., Blackmon, M. H., & Polson, P. G. "A Comprehension-based Model of Web Navigation and Its Application to Web Usability Analysis". In S. McDonald, Y. Waern, & G.Cockton (Eds.), *People and Computers XIV - Usability or Else!* *Proc. of HCI 2000*, (pp.357-373), New York: Springer, 2000.
- [14] Liu, Y.H., Dantzig, P., Sachs, M., Corey, J.T., Hinnebusch, M.T., Damashek, M., & Cohen J. "Visualizing Document Classification: A Search Aid for the Digital Library," *J. of the American Society for Information Science*, 51(3): 216-227, 2000.
- [15] Marlino, M., Sumner, T., Fulker, D., Manduca, C., & Mogk, D. "The Digital Library for Earth System Education: Building Community, Building the Library. *Comm. of the ACM* 44(5): 80-81, 2001.
- [16] Mason, J., Mitchell, S., Mooney, M., Reasoner, L., & Rodriguez, C. "INFOMINE: Promising Directions in Virtual Library Development." *First Monday* 5(6), 2000.
- [17] Mathes, A. "Folksonomies - Cooperative Classification and Communication Through Shared Metadata." Technical Report, University of Illinois Urbana-Champaign, 2004.
- [18] McKay, D. & Cunningham, S.J. "Browsing a Digital Library: A New Approach for the New Zealand Digital Library". *Digital Libraries: Tech. & Mgmt. of Indigenous Knowledge for Global Access. LNCS2911*: 329-339, 2003.
- [19] Nielsen, J. "Search and You May Find." *Alertbox*. [<http://useit.com/alertbox/9707b.html>] 1997.
- [20] Nobel, P. & Shiffrin, R. "Retrieval Processes in Recognition and Cued Recall." *J. Exp. Psychology: Learning, Memory & Cognition* 27(2):384-413, 2001.
- [21] Patrick, T.B., Springer, G.K., Mitchell, J.A., & Sievert M.E. "Virtual Shelves in a Digital Library - a Framework for Access to Networked Information Sources," *J. American Medical Informatics Association* 2(6):383-390, 1995.
- [22] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. "GroupLens: An Open Architecture for Collaborative Filtering of Netnews." *CSCW '94: Proc. ACM Conference on Computer Supported Cooperative Work*, (pp. 175-186), Chapel Hill, 1994.
- [23] Rosenfeld L. & Morville P. *Information Architecture for the World Wide Web*. Sebastopol, CA: O'Reilly and Associates, 1998.
- [24] Spink, A. "Web Search: Emerging Patterns." *Library Trends*, 52(2): 299-306, 2003.
- [25] Sweller, J. "Cognitive Load Theory, Learning Difficulty and Instructional Design". *Learning and Instruction*, 4, 295-312,1994.