

Searching with Tags: Do Tags Help Users Find Things?[†]

Margaret E. I. Kipp* and D. Grant Campbell**

* School of Information Studies, University of Wisconsin-Milwaukee,
Bolton Hall Rm: 510, 3210 N Maryland Ave, Milwaukee, WI, USA 53211 <kipp@uwm.edu>

** Associate Professor, Faculty of Information and Media Studies, University of Western Ontario, London, ON, Canada N6A 5B7 519-661-2111 ext.88483 <gcampbel@uwo.ca>

Margaret E. I. Kipp is an Assistant Professor and member of the Information Organization Research Group, School of Information Studies, University of Wisconsin-Milwaukee. She has a background in computer science and worked as a programmer/analyst. She has a PhD in Library and Information Science from the University of Western Ontario. Her research interests include social tagging, information organisation on the web, classification systems, information retrieval, collaborative web technologies and the creation and visualisation of structures in information organisation systems.



D. Grant Campbell is an Associate Professor in the Faculty of Information and Media Studies at the University of Western Ontario. His research interests include bibliographic description, information and literary theory, metadata and the Semantic Web. He has published in the *Journal of Academic Librarianship*, *Cataloging & Classification Quarterly*, *Journal of Internet Cataloging*, *Epilogue*, and *Papers of the Bibliographical Society of Canada*.



† A preliminary version of this study was published as: Kipp, Margaret E.I. 2008. Searching with tags: do tags help users find things? In Arsenault, Clément and Tennis, Joseph T. eds. *Culture and identity in knowledge organization: proceedings of the Tenth International ISKO Conference 5-8 August 2008 Montréal, Canada*. Advances in knowledge organization vol. 11. Würzburg: Ergon Verlag, pp. 320-35.

Kipp, Margaret E. I. and Campbell, D. Grant. Searching with Tags: Do Tags Help Users Find Things? *Knowledge Organization*, 37(4), 239-255. 28 references.

ABSTRACT: The question of whether tags can be useful in the process of information retrieval was examined in this pilot study. Many tags are subject related and could work well as index terms or entry vocabulary; however, folksonomies also include relationships that are traditionally not included in controlled vocabularies including affective or time and task related tags and the user name of the tagger. Participants searched a social bookmarking tool, specialising in academic articles (CiteULike), and an online journal database (Pubmed) for articles relevant to a given information request. Screen capture software was used to collect participant actions and a semi-structured interview asked them to describe their search process. Preliminary results showed that participants did use tags in their search process, as a guide to searching and as hyperlinks to potentially useful articles. However, participants also used controlled vocabularies in the journal database to locate useful search terms and links to related articles supplied by Pubmed. Additionally, participants reported using user names of taggers and group names to help select resources by relevance. The inclusion of subjective and social information from the taggers is very different from the traditional objectivity of indexing and was reported as an asset by a number of participants. This study suggests that while users value social and subjective factors when searching, they also find utility in objective factors such as subject headings. Most importantly, users are interested in the ability of systems to connect them with related articles whether via subject access or other means.

1.0 Introduction

In traditional subject access systems, the indexer is an intermediary: an individual trained in the rules of information organisation to assign important information about the physical media and the subject matter of the content. On the web, the indexer has typically been the creator of the item, or an automated system collecting basic word frequency information to determine approximate topics. More recently, there has been a growing move to classify materials manually using consensus classifications created on the web by large groups of users tagging material on social bookmarking sites.

Information retrieval research traditionally has been concerned with the efficiency with which information systems retrieve information that is relevant and useful, concerning itself with matters of precision, recall, and system effectiveness. Such studies contain an implicit evaluation of the categorisation of the material (since this affects retrieval) but do not often make this implicit (Cleverdon 1967). This pilot study aims to explore questions pertaining to resource discovery in a new context, that of social tagging. Proponents of tagging and social bookmarking often suggest that tags could provide at worst an adjunct to traditional classification systems and at best a complete replacement for such systems (Shirky 2005). The user-created nature of these organisational schemes suggests that tagging systems may be able to function as a new method for resolving the gap between a user's information need and its translation into a search query by increasing the user's involvement in the categorisation process and combining it with elements of personal information management.

The ability to discover useful resources is of increasing importance where web searches return 300,000 (or more) sites of unknown relevance and is equally important in the realm of digital libraries and article databases. The question of the ability to locate information is an old one and led directly to the creation of cataloguing and classification systems for the organisation of knowledge. However, such systems have not proven to be truly scalable when dealing with digital information and especially information on the web. Can the user-created categories and classification schemes of tagging be used to enhance search in these new environments? Much speculation has been advanced on the subject but so far no studies have examined user perceptions of the utility of tags in a mediated search process.

Social bookmarking tools allow users to store their favourite bookmarks in a publicly accessible manner on the web. Users are encouraged to add descriptive terms or tags to each bookmark. Tagging is the process of assigning a label (whether classificatory or otherwise) to an item and is often combined with social bookmarking or the organisation of other information on the web, for example organising pictures on Flickr.com (Hammond et al. 2005). While other groups have been involved in creating index terms (for example, journal article authors who are asked to provide keywords with their submitted articles), these keywords generally have a small circulation and are not widely used (see Kipp 2005). Small-scale indexing is common but generally covers a narrow range of topics and is specific to the article. Collaborative tagging systems such as CiteULike (<http://www.citeulike.org>) or Connotea (<http://www.connotea.org>) allow users to participate in the classification of journal articles by encouraging them to assign useful labels to the articles they bookmark.

With traditional indexing systems and tagging beginning to coexist, this raises the question: what is the relationship between tagging and traditional indexing systems? Could tags provide a more interactive, mutually-determining relationship, when combined with traditional subject access, that could evolve over time? Or, have systems that have begun to include tags incorporated nothing more than a fad which could lower user expectations of retrieval using traditional indexing systems without providing a similar or better retrieval or indexing performance? While users have assigned many tags to items in social bookmarking systems, there has been little research into how well these tags serve in their suggested function of helping people to re-find the items they had previously located or to enable others to find these items through the use of meaningful tags.

2.0 Related Studies

Previous research in classification suggests that there is a distinct difference between user-created or naive classification systems, on the one hand, and those created by professional indexers on the other (Beghtol 2003). While both systems employ subject-based terms, users tend to employ terms that remind them of current or past projects and tasks, and terms which could have little meaning to those outside their circle of friends and acquaintances, but are very meaningful to the user (Malone 1983; Kwasnik 1991; Jones et al. 2005).

End-user and search thesauri using user-centred and user-generated terminology were developed in the 1980s (Nielsen 2004, 60) to enable users to expand their searches and make connections to thesaurus vocabulary while searching, but many systems still do not offer thesaurus enhanced search (Nielsen 2004, 60). Scholars have also examined usability and user perceptions of thesaurus enhanced search tools and found that these tools enhance the search process, but research into user interactions with such systems is limited (Shiri and Revie 2005; Blocks, Cunliffe, and Tudhope 2006; Shiri and Revie 2006).

Mathes proposes that librarians embrace user assigned tags as a third alternative to traditional library classifications and author-assigned keywords (Mathes 2004), a suggestion which builds on earlier work in end-user and search thesauri. He and others also suggest that user tagging systems would allow librarians to see what vocabulary users actually use to describe concepts and that this could then be incorporated into the system as entry vocabulary to the standard thesaurus subject headings (Mathes 2004; Hammond et al. 2005). Preliminary research has been undertaken in the area of using tagging to generate user centred terms for a thesaurus (Schwartz 2008; Yoon 2009) building on this earlier work with search thesauri. Some libraries and museums have developed systems that attempt to combine the benefits of professional classifications with those of naive classifications by adding tagging to their existing systems. The Steve museum project (Trant 2006), the University of Pennsylvania PennTags project (Allen and Winkler 2007), and Facetag (Quintarelli, Resmini, and Rosati 2006) are all examples of this phenomenon.

Studies comparing the terminology used in tagging journal articles to indexer-assigned controlled vocabulary terms suggest that many tags are subject related and could work well as index terms or entry vocabulary (Kipp 2005; Hammond et al. 2005; Kipp and Campbell 2006); however, the world of folksonomies includes relationships that would never appear in a library classification or thesaurus, including time and task related tags, affective tags, and the user name of the tagger (Kipp 2005; Kipp and Campbell 2006; Kipp 2007). These short term and highly specific tags suggest important differences between user tagging systems and author or intermediary classification systems which must be considered.

Although users searching online catalogues and databases often express admiration for the idea of controlled vocabularies and knowledge organisation systems, they may find it difficult to accommodate

their vocabulary to the thesaurus and often find the process of searching frustrating (Fast and Campbell 2004). Users also tend not to perform the sort of systematic search process common to expert searchers thus limiting their ability to gain the necessary experience with the controlled vocabulary of a system (Markey 2007). Additionally, controlled vocabulary indexing has proven costly and has not proven to be truly scalable when dealing with digital information, especially information on the web (Shirky 2005). Can the user-created categories and classification schemes of tagging be used to enhance resource discovery in these new environments? Much speculation has been advanced on the subject but so far few empirical studies have been done. Heymann, Koutrika, and Garcia-Molina (2008) analyse tags with respect to the pages to which they are assigned. Their research finds that in over 50% of cases, the tags appear in the text of the pages to which they have been assigned. In fact, in 80% of cases, the tags appear somewhere in the text of the page or in the backlink or forward link text from which they were located. They suggest that this positive result means that tags will indeed be a potential asset to improving search (Heymann, Koutrika, and Garcia-Molina 2008), but do users actually use tags when they are present?

3.0 Research Questions

The following exploratory study offers a comparison of the usefulness of a social bookmarking tool and of a traditional online database in an exercise of mediated resource discovery through keyword search. It seeks preliminary answers to the following research questions:

1. Do tags appear to enhance the subjective experience of resource discovery? Do users feel that they have found what they are looking for?
2. How do apprentice librarians find searching social bookmarking sites compared to searching more classically organised sites? How do tags work when searchers are undergoing a learning process with a problem that is not necessarily familiar?
3. Do tagging structures appear to facilitate resource discovery? How does this compare to traditional structures of supporting resource discovery?

4.0 Methodology

Exploratory studies of emerging social phenomena are particularly amenable to qualitative inquiry, thus qual-

tative techniques were employed in the present study. A total of 10 participants were recruited for this study. These participants were recruited from current and former students in library and information science. Current and former students in library and information science were recruited for the following reasons:

1. They may be recent graduates from undergraduate programs, and have retained a memory of their information use in an academic context, or they may have worked for years in an information related field;
2. They have an interest in information issues, which makes them familiar with many online search tools that are popular within the broader online community;
3. As librarians or information scientists, they have become exposed to the vocabulary used to articulate problems that are typically encountered in broader user populations and to empathise with typical user problems in information searching.

Participants were encouraged to compare their experiences with the on-line database and social bookmarking site to their experiences using web search engines in order to increase the volume of data collected about how users select keywords for search. They were also encouraged to talk about their search experiences in the study in relation to past search experiences.

While the use of information science students for this study may suggest a potential bias in the results, there is no reason to assume that all information science students are particularly well versed in the phenomenon of tagging, and there is greater reason for assuming that participants with some experience searching would be able to make the transition between search systems with minimal training, thus removing some of the issues involved with differing interfaces. Library and information science students are expected to learn and become comfortable with a variety of different search systems with varying interfaces. Students in an LIS programme are typically exposed to a variety of search interfaces as part of their education, as opposed to working professionals who may have grown used to a small suite of frequently-used tools on the job.

There have been no empirical studies on the experience of users using tagging systems in an LIS context. Given the increasing interest in such qualitative data as user relevance judgements (Tang and Sun 2003; Oppenheim, Morris, and McKnight 2000),

this study will examine the qualitative dimension that shows how controlled vocabularies, user index terms and tags relate to each other. Because of the emphasis on the qualitative dimensions of this exploratory study, the study is limited to a small number of participants. The results of the study involved the triangulation of three primary data sources: interviews, search terms and screen captures of search sessions.

The researchers were asked to search PubMed (an electronic journal database of articles for use by researchers and practitioners in the health sciences) and CiteULike (a social bookmarking site specialised for academics with a wide range of health sciences articles already tagged by users) for information on a specific assigned topic (see Table 1). The topic was provided as a paragraph describing an information need:

You are a reference librarian in a science library. A patron approaches the reference desk and asks for information about the application of knowledge management or information organisation techniques in the realm of health information. The patron is looking for five articles discussing health information management and is especially interested in case studies, but will accept more theoretical articles as well.

This topic was chosen by the researcher after searches showed that there were sufficient articles on the subject of information management techniques used in health information in both databases that participants would be able to find far more than the number of relevant articles requested.

Screen capture software (specifically CamStudio and Xvidcap), a "think aloud" protocol (Krug 2006), and a semi-structured exit interview were used to capture the impressions of the users when faced with traditional classification or user tags and their usefulness in the search process.

Each participant searched for information using both the traditional on-line database with assigned descriptors and a social bookmarking site. Participants were asked to perform the searches in the order specified so that their use of a social bookmarking site first versus an on-line database could be alternated to compensate for order effects.

Participants selected their own keywords for searches on both tools after having read the paragraph description of the information need. They were then asked to provide a list of terms they would use to

Activity	Description	Length
Welcome	initial greeting and welcome	2-3 minutes
Introduction to session	Introduction to the study discussing the session itself and the tasks they will be asked to perform.	5-7 minutes
First search task (CiteULike or PubMed)	The first of two tasks consisting of: 1) the user's generation of keywords for search, 2) collection of articles, 3) analysis of retrieved articles for relevance, and 4) assignment of relevance judgements to the articles, 5) assignment of new set of keywords for search	15 minutes
Second search task (PubMed or CiteULike)	same as first task	15 minutes
Post search discussion	A semi-structured interview involving a discussion of the participant's results and their own thoughts as to the usefulness of the terms they used to search and the terms used to describe the documents they retrieved.	15 minutes
Conclusion	Final comments and a thank you for participating.	3-5 minutes

Table 1. Preliminary Timeline for Sessions

start their search. Participants were asked to search until they had located approximately five articles that appeared to match the query and assign relevance scores to articles based on an examination of available metadata. At the end of the search process, participants were asked to make a second list of terms they would now use if asked to search for this information again. Participants did not have access to their initial set of search terms at this time to eliminate the learning effect. Participants' actions were recorded using screen capture software and a microphone. Additionally, participants were interviewed after the search process in order to allow them to articulate their impressions of the search process.

The following questions were used as a guide in the semi-structured interview:

1. Did you find the user assigned tags were a better match for the keywords you chose initially? If not, were they useful in locating the relevant articles? (Also ask this question with respect to subject headings.)
2. Did you find the subject headings useful? Would you have used any of the subject headings or tags to index the document? Would you use any of the subject headings or tags to search for this document again?
3. Now that you have performed the search, what do you think of the differences/similarities between your initial and final sets of keywords? (Depending on the responses, it may also be useful to discuss individual keywords, especially keywords

that may have been dropped from the search process or that were dropped during the search process only to reappear in the participant's final list.)

4. What are your thoughts on keywords or tags which you chose not to use in your search?

One issue that might have had an effect on data collection is that of differing user interfaces; however, both CiteULike and PubMed offer search by keyword and participants were given a brief introduction to searching with both systems (including an introduction to the MeSH browser in PubMed and the tags in CiteULike). Participants with a library and information science background were specifically chosen for the study because of prior experience with searching multiple systems with different interfaces, so that they would be better able to handle differences in interfaces. The design of this study is based on common information retrieval research designs with an emphasis on the collection of keywords used in the search (as in web log analysis) in addition to the collection of a ranked set of documents judged relevant by the participant.

Three sets of data were thus available for analysis: sets of initial and final keywords selected by the user, the recording of the search session and think aloud, and recorded exit interviews after the search session. These three data sets were examined to balance the users' perceptions of the search (interviews) with their search strategies (terms) and their behaviour while implementing those strategies (screen captures). Keywords and tags chosen by users were

compared and examined to see how or whether they were related and participant's recorded video sessions were transcribed along with the interviews in order to provide a deep analysis of the search process of the study participants. These transcripts were then analysed using a grounded theory approach (Strauss and Corbin 1990) based on initial insights while transcribing the video sessions, beginning early in the observation process. This coding was then used to aid in choosing what search behaviours to look for in the transcripts. Trustworthiness of the results was ensured through a triangulation of participant experiences, deep analysis of the results, and discussion between the researchers.

5.0 Results

5.1 Demographics

A total of 10 participants were recruited for this study. Four of the participants were male and six were female. Participants were between 23 and 40 years of age and generally self-identified as intermediate level computer users (80%) while the remaining participants (20%) self-identified as expert users. All but one of the participants listed previous educational backgrounds in the humanities (English and French) or social sciences (Political Science, Sociology, etc.). The final participant gave an educational background in the fields of mathematics and education. Professional backgrounds were generally in the areas of teaching or librarianship/archives; however, 3 of the participants did not include a professional background.

Number of years using a computer ranged from 6 years to 22, with a median of 19 years of experience using a computer. Participants were chosen from amongst users who have some experience searching the Internet, so it is reasonable that all participants would have some experience with computers. Participants' use of specific Internet tools was mixed. Only 20% of participants reported having a website, and 40% a blog. However, one of the users with a blog also maintained a webpage. Half the participants maintained neither a blog nor a website. Participants were generally frequent users of both web search engines and journal databases, and therefore were reasonably conversant with both searching and web use; but, they were relative novices at tagging systems. Ninety percent (90%) of participants used search engines often or frequently and 70% of participants used journal databases often or frequently.

While participant use of search engines and journal databases was high, few participants reported using social bookmarking tools on a regular basis. Fully 70% of participants reported using them rarely or never. Social bookmarking tools are still relatively new, especially in comparison to journal databases, and heavy users are still less common.

5.2 Participant Keyword Usage

All users used multi word keywords initially, suggesting that the users are indeed experienced searchers who are aware of methods which can be used to improve precision or recall in search. At the end of the search process, when users were asked to generate a new list of keywords they would now use for the search, half the users separated their list of final keywords by tool, despite the fact that they were asked for only one list.

A total of 28 unique keywords or keyword phrases were listed initially by the participants. These keywords and keyword phrases were entered into the system by participants according to the patterns discussed later in this paper. Each participant listed between 1 and 9 keywords initially, with the median value being 6 keywords.

Keyword	Frequency
knowledge management	7
information organisation/information organization	6
health information	6
case studies/case study/“case stud”	4
health information management/health info mgt	3

Table 2. Initial Keywords

The four most commonly chosen terms were: knowledge management, information organisation, health information, and case studies (Table 2). Each of these terms is directly from the initial text of the information need. Users reported that their use of knowledge management versus information organisation during the search process was determined by the types of results they found when searching with each tool. The fifth search phrase is a reasonably obvious contraction of health information and knowledge management.

Participants produced 46 unique keywords for their final lists (Table 3). They used between 3 and

16 keywords in their final lists, with the median being 6. Participants who separated their final lists by tool used between 3 and 8 terms for CiteULike (median 5) and between 1 and 8 for PubMed (median 3). One participant chose the term "Information Management" which is a MeSH descriptor as the only keyword for searching PubMed.

Keywords	Frequency
knowledge management/km	9
case studies/case study	6
health information	5
information management	5
health care	3
health information management	2
informatics	2
health	2

Table 3. Final Keywords

The most commonly used keyword, by far, was knowledge management. This term comes directly from the information need as described above and is in keeping with previous information retrieval studies where users tended to select terms from the given text of information need for search (Oppenheim, Morris, and McKnight 2000). Information management was also a commonly-used term; this term could be seen as a modification of knowledge management to fit the terminology of a different group of users who prefer the term information management. Another commonly-chosen term was health information, also from the information need. Both information management and health information were tied for third most popular for the two tools. While users often mentioned that they considered their initial keyword sets to have been incomplete, they tended to choose the same or very similar terms as their suggestions for good search terms to use in order to produce better results. This suggests that their initial search terms were well chosen and matched closely those chosen by users tagging articles in CiteULike, but also came close enough to terms used in the Medical Subject headings used in PubMed (or its entry vocabulary) or terms used by authors whose works are published in PubMed for good results to be retrieved.

Half the participants separated their final keywords lists by tool (Table 4). Again, knowledge management was the clear favourite, having been chosen 6 times in total and 4 times for CiteULike. Opinion

was more split on whether knowledge management or information management were best for PubMed. Participants who discovered that information management was a MeSH descriptor were more likely to suggest this as the preferred term while other participants found that knowledge management was useful for free text searching of abstracts.

Keywords	CiteULike	PubMed
knowledge management	4	2
information management	1	3
case studies	3	1

Table 4. Most common terms separated by tools.

Case studies is not a descriptor in PubMed, but it is an entry term for the descriptor "case reports" that includes case studies. Since this term is an entry term for a MeSH descriptor, it will allow the user to connect directly to the MeSH vocabulary without having to search for a specific term as was the case with information management.

The other popular term, knowledge management, is not a descriptor or an entry term in MeSH, but it can be used to retrieve articles through free text searching of abstracts. Knowledge management was not as frequently chosen for use in PubMed because many participants found that it was not as useful a search term since it is not a MeSH descriptor. Knowledge management and information management are very similar concepts since they both deal with the organisation of information into a form usable by others, but the terms tend to be used in different fields. The high use of knowledge management in this study and on CiteULike suggests that MeSH would be well advised to consider how the term would fit into their descriptors as an entry term, at minimum.

In all, participants suggested 20 unique terms for use in searching CiteULike (18 were used by only one person) and 17 unique terms for use in searching PubMed (15 were used by only one person). This wide spread of suggested terms used by only one person is additional evidence for the existence of the long tail in tagging and searching and supports studies showing that searchers do not use the same terminology when tagging (Kipp 2005; Kipp 2007).

5.3 Participant Search Experiences

Participants tended to prefer the search experience on the system used first, regardless of previous experi-

ence with either system or similar systems. Further interviews may be required to determine whether this trend continues although it might simply be the case that any frustration with the system used second would still have been uppermost in the participant's mind.

"PubMed just didn't seem as useful. Though I don't know whether these articles [CiteULike articles] are going to be as academic as something in PubMed. If they're from the core journals or not." – **Participant 1** (used CiteULike first and had prior experience with Ovid and Medline, but not PubMed)

In contrast to participant 1, participant 9 did not like the CiteULike interface and was much more impressed with the PubMed interface and its features, but "would have liked to have subject headings visible along with [the] abstract." Participant 10 explicitly stated that the PubMed search was easier than the CiteULike search and that CiteULike's lack of an advanced search box and a search history made it much less useful.

Other participants found that the interface was providing too much rather than too little information. Participant 1 felt that the PubMed interface was overwhelming and preferred a simpler interface with slightly less information upfront.

"I think if I knew how to use PubMed better I might have been able to get better results but I don't have the experience. It was just a little overwhelming. Too many results. ... Like in a Google search. ... I can't really tell how many results I was finding with CiteULike. I did find it useful in PubMed how they linked to related articles. That was useful." – **Participant 1**

Participants expressed frustration with the interface and the use of keywords in the systems. In general, participants expressed the impression that their use of both systems was hindered by the problems of learning different and complex interfaces including: the locations of search boxes, identification of controlled vocabulary terms, different sets of metadata displayed in the results, and other features of each system.

"I found it a lot easier to search CiteULike for some reason. I'm not sure. I think with PubMed I could find some better keywords, keywords

that might be indexed. It looked like with CiteULike I could just type in things like health care, health organisation." – **Participant 1**

Participant 7 expressed a similar view and stated that the PubMed search was frustrating because it was difficult to figure out which terms to use. In contrast, Participant 2 explicitly stated a preference for Google after the search process. The participant described significant search experience on Google and felt that this experience did not translate directly despite the familiar interface of the search box.

"I found that it was sort of frustrating because I wasn't familiar with the databases. If I had been more familiar, if I had more experience, maybe I would have been able to narrow the keywords faster. Um, yeah, that was it and also being limited to those two databases, um, I would have tried Google. I love Google. I just go onto Google and then what I would do is I would—when I do information searches it's more scatter brained. I would find one article and I might read through it and then it might suggest something in the article that would lead me to another source and I would look at that and... so it's more of a, um, following the breadcrumbs sort of way to do things." – **Participant 2** (participant describes favoured citation pearl growing search strategy on Google)

This is an interesting finding because many search systems seem to be explicitly assuming that users will be comfortable with basic searches since Internet searching is so common. This comment, however, suggests that users may be assuming that there is considerable complexity in other search systems that they do not understand and therefore are unable to access. Additionally, these users appear to be concerned that this complexity is keeping them from making full use of the system, this despite the fact that Google's organisation is equally complex and it is almost impossible to be sure one is making full use of Google.

"I really should have looked more closely into how their [CiteULike's] search function worked, because I know I included health, but I'm not sure if it's assuming the AND operator. So I was getting a lot of stuff that was on knowledge management but not necessarily anything to do with health." – **Participant 6**

The most popular form of metadata as articulated by the majority of participants in the post search interviews was the abstract. Participants frequently lingered over abstracts and occasionally complained aloud during the search process if the abstract was missing.

Interviewer: "I'm interested in what metadata people find useful when searching. If the lack of an abstract is a huge deal..."

Participant: "It is a huge deal. You can't tell anything about the article without it." – **Participant 1**

While participants listed the abstract as the most important piece of information for determining relevance, they also stated that titles or links to related articles were just as useful as, or even more useful than, subject headings or tags.

"I mostly just looked at the titles of the article, read a little bit of the abstract and then the keyword that I used. I would give that to the user and it would be up to them to decide if the articles were in fact useful and they could continue the search from there. ... I did find it useful in PubMed how they linked to related articles. That was useful." – **Participant 1**

In fact, many participants felt that the tags were most useful as links to related items rather than as guides to subjects. One participant claimed not to have used the tags, but found the related articles listed in PubMed very useful. This participant thought that if asked to repeat the search again that the tags would be useful as a form of related article search.

"[I thought] I wasn't using the tags, but I was actually using them to look at related articles" – **Participant 10**

"It [the tags] might have been useful for searching but like I was looking for specific things like case studies into information management in health care and uh in order to know if the article was relevant or not I had to go into the abstract and you know if the abstract seemed, um, relevant than I would look into the full article you know to get a better idea of whether it's good or not." – **Participant 2**

Participant 9 reported that it would have been helpful to be able to "select combinations of tags by

clicking on them" a feature which has recently been implemented on another social tagging service, Del.icio.us. This would be similar to the PubMed feature whereby users can combine previous searches to create a new search.

In addition to title, author and abstracts, participants also made use of keywords in PubMed. Some participants made use of various features of PubMed including the details tab which displays their query modified with automatically chosen MeSH headings where appropriate and the MeSH browser itself to select useful keywords for search.

Many participants found that searching PubMed fit with their previous search experience searching journal databases and were quite comfortable with this part of the search process. Both participants 4 and 6 stated that the PubMed interface was much more friendly since it provided a typical online database searching experience with a thesaurus while CiteULike had only user tags. Participant 10 echoed this view, and suggested that the tags were too narrow to be useful as opposed to the MeSH subject headings. Other participants found that their terminology did not match that used in PubMed and that the MeSH browser did not always provide an alternative.

"What I started off with, what I started off with was using some of the words in here [the initial information need] like knowledge management, information organisation and so on. ... And in PubMed when those words didn't work and I was getting nothing, that's when I started branching out and putting library and trying to figure out like different synonyms, synonyms or uh." – **Participant 2**

Participant opinion was also split on the utility of the tags. Many participants felt that the tags were an excellent addition to the system, while others felt they were either too broad or too narrow for an effective search.

"Um, I found that a lot of the keywords I used were already used as keywords in CiteULike, so I think they were good keywords. To use. But because they list several keywords along the bottom, I can pick up new ones as I go. And again, because they're only one word, I can remember them. Public health, ehealth, health services, it was a kind of recurring term on a lot of the articles that I thought would be useful." – **Participant 5**

"Well, I didn't really find these tags to be particularly useful to be honest. One of the things that kind of bothered me about them is that they weren't really grouped... you have care and health but you don't have health care together. You have care, health and informatics. It would be useful if it was healthcare and informatics together as one tag. Instead, because if you just click on health. It's not applicable at all, you know, and like km is a term, but then knowledge and management are separate, which is kind of bothersome." – **Participant 1**

Participant 1 included the tag "km" in the final list of keywords, despite having found the tags to be problematic when compared to the more familiar controlled vocabularies of traditional databases. Despite not personally deeming the tags useful, the participant must have felt that this tag could be useful to other searchers using CiteULike. A number of participants commented on the use of different terminology for different systems and as previously noted many insisted on dividing their final keyword lists by tool.

"Hmmm. Because this is PubMed, we probably don't need health in here. Because everything is health. Okay, and I probably wouldn't use km either. It might not be as, uh, common in PubMed" – **Participant 2**

Some participants expressed some confusion at the differences in the visible organisational structures used by PubMed and CiteULike. These participants showed or discussed their confusion when faced with the differences between keywords and tags and the methods used to organise and retrieve information in the two different online databases.

Interviewer: "OK, now. Which one did you like the best?"

Participant: "Oh, the first one, CiteULike."

Interviewer: "What did you like about it?"

Participant: "Just because there was more words, reference words. After the words I put in.... they just eventually appeared. I don't know what I was doing." – **Participant 8**

This result suggests that even library and information science students can suffer from confusion when faced with a new and unfamiliar system. Systems where the organisational structures are hidden from them, such as Google, conversely seem to offer less confusion

since users do not seem to feel they need to know anything about how the system works. This may be due to the fact that Google is almost certain to return something no matter how little knowledge a user has of a subject (Fast and Campbell 2004). As participant 7 stated, "It was easy to kind of, uh, expand my search by just clicking on tags. I felt like on PubMed I had to find that one, uh, word that they used."

Some participants confused tags and descriptors or expressed an unfamiliarity with the concept of multiword subject headings. Participant 5 expressed such concerns stating that the tags on CiteULike were more friendly because they were shorter, ignoring that many CiteULike tags are in fact multiword tags joined by various punctuation marks.

"Oddly enough, CiteULike, which is totally regulated by users, I actually found to be the most similar to Library of Congress: again it picks one, short, nice, concise words as subject headings, that lead into a nice broad topic that I can move around in and play with. Um, PubMed was a little unlike anything I'm used to. Its descriptors were just too long. I'm sure I could make a go of it eventually, but just sitting down to try initially, it is a little more work than it should be. Even things like digg and delicious, the keywords are usually 2 words long, maybe three. And that actually might be why I find CiteULike easier to use; it's similar to what I'm used to, like dig and delicious." – **Participant 5**

A number of participants discussed issues with the interfaces of each system and specifically with the organisational systems used in each system. As previously noted, Participant 9 felt that CiteULike should support the ability to quickly combine tags by clicking on them, a form of filtering for results which is present in some journal databases and library catalogues (e.g., Endeca <http://www.endeca.com/>), Endeca's ILS system allows faceted browsing and filtering.

Other participants expressed a desire for more order in online systems, despite often having expressed confusion when faced with this order. This juxtaposition of a user-defined need for order and a user-expressed confusion when faced with structured and controlled vocabularies poses significant issues for system designers.

"It would be nice if there was a coherent structure to it as opposed to the way they've [CiteULike] done it here. Um, other thoughts, I

think if I knew how to use PubMed better I might have been able to get better results but I don't have the experience. It was just a little overwhelming."— **Participant 1**

Participants suggested that CiteULike should adopt additional information organisation techniques and did not in general mention tag clouds or tag lists as options. Despite this, participants also occasionally expressed frustration with PubMed's search and suggested that subject headings should be more prominently displayed in the search results. "[I] wanted to be able to have subject headings [in PubMed] visible along with the abstract."— **Participant 9**

Participants also noted that in addition to tags, CiteULike also offers the feature that you can see who posted the article and then see other articles and other tags by this same user. "You can search by tags or you can search by people and it also shows the people who are interested in this idea... this search term that I put in."— **Participant 7**

This ability to see another person's tags and articles is a feature that does not have an analogue in a traditional journal database. While tagging itself is similar to the use of controlled vocabulary headings, the association of a user or group with a set of articles is not normally present in a system and such associations are made much more haphazardly by, for example, a colleague's email about an article. Often, participants seemed to be searching for recommendations, a personal touch, in the tags. They appeared to be figuring out that once they were in the right subject area, the tags applied by a particular user could be helpful to them and serve as an important guide to the relevance of tagged items.

While participants' views were solicited on the search process and their use of interface features and keywords, a key component of this study was the examination of the differences between participant keyword use, statements made in interviews, and the actual search behaviour of participants. While participants were often quite articulate about their search preferences and behaviours, some inconsistencies were observed between participant's expressed preferences and actual behaviours.

5.4 Participant Search Behaviour

When searching, most participants started with a single keyword or keyword phrase, but quickly added additional keywords from their initial lists in order to reduce the number of results returned. Some par-

ticipants immediately made quick assessments and modifications to their initial queries, while others took more time to scan the results. Most participants showed a preference for one or the other behaviour but did show some willingness to change behaviours slightly during the search depending on the number of results.

Keywords: health km case studies

Actions: scrolls slowly down then up again

Keywords: knowledge management case studies

Actions: scrolls more rapidly down the page then up again

Keywords: information case studies

Actions: scrolls part way down then up again

Keywords: library case studies

Article: Realizing what's essential : a case study on integrating electronic journal management into a print-centric technical services department (PubMed: 17443247), does not select

Keywords: "information management"

— **Participant 2**

Many participants showed evidence of uncertainty or frustration when searching one or the other system. Participants paused for longer periods, scrolled up and down without making a selection or hovered over items without selecting anything. Many participants also appeared to be browsing the results on the first page to see if they were getting enough relevant results from their search terms before narrowing or broadening their search.

examines metadata, hovers over journal name, hovers over author name, does not select

— **Participant 9**

Pauses for quite some time before scrolling up and down the hit list. Doesn't go past p. 1

— **Participant 5**

Public health information

doesn't scroll: just clears search box

Education and health care

no scrolling; clears search box again

— **Participant 8**

Participants seemed to occasionally be confused by the differences between controlled vocabularies (such as MeSH descriptors) and tags. It was fairly common for participants to use incorrect terminology to identify their use of terms when searching.

"Um, yes. I found it difficult to actually determine what's relevant, because the subject headings—they're basically a sentence. And remembering what's been said, if there's 1 2 3 4 5 in each one and I have 2 or 3 up, its kinda hard to determine a pattern. ... They did, in that, um if I could remember any recurring words in those sentence-long subject headings, I could write them down and try them again for the next search. It wasn't as easy as remembering one key word on CiteULike, it was trying to read a sentence, picking what might be an appropriate term from that sentence, read the next sentence, and try to compare the two sentences for matching key words that might be useful. It was a lot more work, PubMed..."

– Participant 5

health km case studies

scrolls slowly down then up again

"Hmmm.. Because this is PubMed, we probably don't need health in here. Because everything is health. Okay, and I probably wouldn't use km either. It might not be as, uh, common in PubMed so..."

– Participant 2 (initial search on PubMed)

All participants used Boolean searching in both PubMed and CiteULike in order to narrow their search and appeared to expect it to be present as only a few of the participants asked the interviewers if Boolean search was supported. Most participants also used truncation, again expecting it to be supported. One participant even used the near operator in a search of CiteULike. Like PubMed, CiteULike does indeed support truncation, wild cards and Boolean search (though only with symbols) but it does not in fact support near as an operator (http://www.citeulike.org/search_help).

"information 2N organization" and "health information" and "case stud**" – Participant 10

All participants used internet searching techniques such as quotations to indicate a phrase search and many also dropped the AND in boolean searches as expected on Google.

Many participants expressed a desire for an abstract with the retrieved records on PubMed and CiteULike and their searching behaviour bore out this desire. Participants selected, hovered over or scrolled slowly through abstracts and even parts of articles to determine relevance.

user examined article 561415, scrolled past other metadata to read abstract – Participant 2

scrolls up and down, locates article link and selects, scrolls to read first few pages of article – Participant 2

Tags were used by a number of the participants despite many claims to the contrary. However, participants may not have felt that their use was sufficiently close to the concept of "using a tag as a search term" to constitute the sort of use the interviewers wanted. A number of participants stated that they did not use the tags, although they had clicked on or otherwise examined them or even used them in query lists as participant 2 did in the previous excerpt. This suggests that participants may see clicking on subject terms in order to browse the results as a distinct activity from searching using a subject term.

"One of the articles used km. I wonder if that would help." – Participant 2

Query: "health information" km "case stud"

– Participant 2

selects tag labelled healthcare – Participant 10

Scrolls down list and hovers over tags momentarily – Participant 3

mouse hovers over tags; clicks on tag bioinformatics – Participant 9

Selects tag "health-information" from first article in hit list

Get's "cyrille's health information [8 articles]"

– Participant 5

clicks on tag partners-in-health, but does not select article, returns to main list

Health information systems: failure, success and improvisation (CiteULike 312350)

pauses over abstract for a short period, then selects this article

clicks on tag health-care, scrolls down, scrolls up and returns to main search list

– Participant 1

Participants also used descriptors in PubMed. Some even selected these descriptors from the MeSH browser or the details tab after an initial search.

“Um, really only 2 that immediately jumped out; um, managed care seems to be actually like the key term for both of them. So, if I were to continue I’d probably search that to see what else comes up.” – **Participant 5**

Actions: Examines details tab (“Health Inf Manag”[Journal] OR “HIM J”[Journal] OR (“health”[All Fields] AND “information”[All Fields] AND “management”[All Fields]) OR “health information management”[All Fields])
 Keyword: Health information management
 – **Participant 3**

selects MeSH search to find keywords
 health information management
 clicks on Management information systems
 – **Participant 9**

Participants used a number of other features of both systems including related articles links in PubMed and group names in CiteULike. This suggests that it would be most useful to provide users with a list of other items with similar subject headings or tags and as much additional metadata as possible to allow the user to browse related items by as many different definitions of related as possible (see Ockerbloom 2006). A related article style feature has been implemented in the University of Pennsylvania’s Online Books page subject search as a test (<http://onlinebooks.library.upenn.edu/subjects.htm>).

Selects an article after a traditional keyword search then returns to main list, scrolls slowly. Returns to previously selected article. Clicks on user name Evidence-based-medicine (group). Scrolls slowly. Selects article: Information retrieval and knowledge discovery utilising a biomedical Semantic Web (CiteULike 405826) – **Participant 9**

Selects tag cloud for user who posted the [current] article. Hovers briefly, selects list of [this user’s] recent articles. – **Participant 4**

A number of participants selected articles from article lists that had been posted under a specific tag by a specific user or user group on CiteULike. While tags themselves can be seen as an analogue to subject headings or descriptors in a traditional journal database, there is no real analogue in traditional information organisation to that of the CiteULike user or group.

This recognition that specific users may provide an additional level of information organisation is a new feature of social tagging systems. Even users who did not actively use user or group names in their search process showed recognition of the presence of users.

“Um, I found this one [CiteULike] easier to navigate, just because of having actual key one-word subjects. So, I’m looking for knowledge management, then I can just type in knowledge management, and if that user’s already bookmarked lots of articles on knowledge management. I can see what they have on their list. Yeah, I found this one much easier to use.” – **Participant 5**

selects tag cloud for user who posted the article; hovers briefly, selects list of recent articles
 – **Participant 9**

Health services (494 articles)
 scroll down
 mouse-over username
 groups interested in health services
 back to search box
 – **Participant 8**

One participant did not find anything useful on CiteULike using the tags by themselves; in fact that participant stated that they were too narrow, but did use user and group names to select articles, finally selecting an article from a user group on CiteULike and an article from a user’s list of articles.

In addition to subject terms such as descriptors and tags, users made use of other special terms for searching, specifically journal names.

Keywords: “Health Inf Manag”[Journal]
 Actions: Scrolls down slowly, selects article
 Article: Health online: a health information action plan for Australia (PMID: 11143002)
 Notes: After selecting this journal, participant selected all other articles from this list by simply scrolling until an interesting article was reached, occasionally, the participant scrolled back up to an article slightly higher on the list
 – **Participant 3**

selects journal name as search term
 J AHMA[Journal] – **Participant 2**

Additionally, participants used the related article links in PubMed to locate relevant articles. Many participants praised this feature and considered it to be just as important or possibly more important than subject headings for locating relevant articles.

Goes to this article; scrolls down (scanning abstract); goes to Related Links; mouseovers different links. – **Participant 5**

“It’s too bad there’s no abstract.” does not select article, but examines related articles on the side and selects one. – **Participant 2**

Action: pauses for a long time over an abstract, decides to select after all but is not sure of relevance, returns to main list and scrolls slowly

Participant: “Would it help to use these related links?”

– **Participant 1**

One participant suggested that the tags were actually most useful as a form of related article link, rather than as subject headings. “[I thought] I wasn’t using the tags, but I was actually using them to look at related articles” – **Participant 10**. This participant showed an awareness of the relationships between the tags assigned to the same article and tags assigned to multiple related articles and was able to suggest a way in which tags or subject headings could be used to enhance traditional search systems by providing explicit lists of articles with similar tags or subject headings rather than just supplying a list of subject terms.

Despite the fact that participants exhibited a fair amount of thought and care in the selection of their keywords and in the use of additional features for locating relevant materials, many participants spent a great deal of time scrolling through long lists of results or entering minor variations on their search query and anxiously examining the size of their result sets.

Notes size of result set and tries another query without scrolling

– **Participant 1**

“That didn’t work.”

Actions: participant continually enters keywords, performs the search and does not scroll before entering new search terms

Keyword: information management

Keyword: information organization

Keyword: knowledge management

Keyword: knowledge management case studies

Keyword: “information science”

Keyword: knowledge organization

– **Participant 2**

These behaviours suggest that users were concerned with selecting good sources and did not find that searching with keywords all by itself was sufficient to help them reach this goal. Many participants praised such features as the related article lists provided in the PubMed interface, and other participants made use of tags and tag clouds, user names, and even group names in CiteULike to help them locate promising relevant articles that were related to an article they found relevant, a set of keywords they felt were relevant or a user who appeared to be collecting relevant articles.

6.0 Discussion

This study examined the relationship between user tags and the process of resource discovery from the perspective of a traditional library reference interview in which the system was used, not by an end user, but by an information intermediary who was trying to find information on another’s behalf. Searching by an intermediary, or mediated search, is a traditional library and information science task tied directly to important library skills in information sources and services and information organisation. Strong LIS elements were present in the search behaviour of the participants. Participants discussed the importance of learning how the search function works on a system when beginning a search and how this can affect the results. They discussed narrowing and broadening searches and selecting specific terms as search terms. They used Boolean search, truncation, and even the NEAR operator. They talked about finding different synonyms and antonyms, and were aware of the common (to librarians) paradox that in a health database the word “health” is so common that it could almost be considered a stopword. Participants were able to bring a set of LIS perspectives to the search process, regardless of their relative skill or lack of skill in searching, which helped to frame their expectations for each system. Although this could be seen as a limitation of the study in terms of application to broader user groups, it provides real insight into how tagging systems could be adopted into library and information science systems and practices.

One issue that cannot be ignored in information retrieval studies is Google. Google’s pervasiveness,

search techniques, assumptions, and interface have become such a large part of the common Internet experience that all search systems are judged against its apparent ease of use (Fast and Campbell 2004). Participants in this study used many Google style search techniques and assumptions, including adding additional keywords from the initial lists, in order to reduce the results returned. In many cases, participants assumed the use of Google style Boolean search where the AND is simply understood as well as the use of quotes to signify a phrase search. All of these search behaviours suggest that Google style search has become a standard, thus perhaps explaining the confusion felt by some participants when using systems with more obviously complex features. If this is true, tagging systems and library systems will need to consider the impact of the confusion caused by the fact that these systems demand more than the ubiquitous Google search box.

7.0 Conclusions

The preliminary study showed that participants did use the tags to aid in the search process, selecting tags to see what articles would be returned. They also used the tags as a guide to suggest further search terms, suggesting that users do indeed pay attention to subject headings and metadata if they fit a pattern users recognise or make sense in the context of their existing knowledge on the subject. Interestingly, many participants stated that they had not used the tags, though examination of the search process showed that they had been using them as links to related articles or sources of search terms. It is possible that they had not considered this to be a full use of the tags as they were not necessarily using the tags as subject headings or search terms.

Participants generally used the same number of keywords for both lists, though many insisted on dividing the final keyword list up by tool. Despite this, the most commonly used terms tended to be the same in each case and knowledge management was generally selected as a useful term for each tool despite the fact that it is not present in MeSH as a descriptor or as entry vocabulary.

Participants reported a number of interface issues which they found degraded or enhanced the search process. Items such as the presence of full metadata, abstracts, and even full text links to articles were lauded, while lack of vocabulary terms and, especially, missing abstracts were deemed to be impediments to search. Participants found related article links and

other newer features of systems to be a significant enhancement to the search process, and some participants reported or were seen using tags or user names in CiteULike for similar purposes.

These findings suggest that users would find direct access to the thesaurus or list of subject headings showing articles indexed with these terms to be a distinct asset in search. Many of the participants in this study made use of the related articles links provided by PubMed and were intrigued by the possibilities of the tags on CiteULike but did not find that the structures were in place to fully support browsing of related items by keyword or combination of keywords.

As shown by Ockerbloom (2006) and in previous research into end-user and search thesauri (Nielsen 2004, 60; Shiri and Revie 2005; Blocks, Cunliffe, and Tudhope 2006; Shiri and Revie 2006) these webs of related items can be built automatically using existing thesaurus structures and displayed to the user. This suggests that indexing and classification structures are fertile ground for the development of newer and better interfaces to document collections as demonstrated by the interest in browsing and combining tags to create a web of related documents, a web which often already exists in traditional databases but has generally been hidden from the user's view.

References

Allen, Laurie, and Michael Winkler. 2007. PennTags: Creating and using an academic social bookmarking tool. *Proceedings of the ACRL 13th National Conference, Baltimore, MD, USA, March 29-April 1, 2007*.

Beghtol, Clare. 2003. Classification for information retrieval and classification for knowledge discovery: Relationships between "professional" and "naive" classifications. *Knowledge organization* 30: 64-73.

Blocks, Dorothee; Cunliffe, Daniel; and Tudhope, Douglas. 2006. A reference model for user-system interaction in thesaurus-based searching. *Journal of the American Society for Information Science and Technology* 57: 1655-65.

Cleverdon, Cyril. 1967. The Cranfield tests on English language devices. *Aslib proceedings* 19n6: 173-194.

Fast, Karl V. and Campbell, D. Grant. 2004. 'I still prefer Google': University student perceptions of searching OPACs and the Web. *Proceedings of the 67th Annual Meeting of the American Society for Information Science and Technology, Providence*,

Rhode Island, USA, November 13-18 Vol. 41, pp. 138-146.

Hammond, Tony; Hannay, Timo; Lund, Ben; and Scott, Joanna. 2005. Social bookmarking tools (I): A general review. *D-Lib magazine* 11n4. Available <http://www.dlib.org/dlib/april05/hammond/04hammond.html> (accessed May 31, 2009).

Heymann, Paul; Koutrika, Georgia; and Garcia-Molina, Hector. 2008. Can social bookmarking improve web search? *First ACM International Conference on Web Search and Data Mining (WSDM'08), February 11-12, 2008, Stanford, CA, USA*. Available <http://ilpubs.stanford.edu:8090/858/> (accessed May 31, 2009).

Jones, William; Phuwanartnurak, Ammy J.; Gill, Rajeep; and Bruce, Harry. 2005. Don't take my folders away! Organizing personal information to get things done. *CHI 2005, April 2-7 2005, Portland, Oregon, USA*. Available <http://hdl.handle.net/1773/2031> (accessed May 31, 2009).

Kipp, Margaret E.I. 2005. Complementary or discrete contexts in on-line indexing: A comparison of user, creator and intermediary keywords. *Canadian journal of information and library science* 29: 419-36. Available <http://dlist.sir.arizona.edu/1533/> (accessed May 31, 2009).

Kipp, Margaret E. I. 2007. @toread and cool: Tagging for time, task and emotion. *Proceedings of the 8th Information Architecture Summit, Las Vegas, USA, March 22-26*. Available <http://dlist.sir.arizona.edu/1947/> (accessed May 31, 2009).

Kipp, Margaret E.I. and Campbell, D. Grant. 2006. Patterns and inconsistencies in collaborative tagging practices: An examination of tagging practices. *Annual General Meeting of the American Society for Information Science and Technology, Austin, TX, USA, November 3-8, 2006*. Available <http://dlist.sir.arizona.edu/1704/> (accessed May 31, 2009).

Krug, Steve. 2006. *Don't make me think: A common sense approach to web usability*. 2nd ed. Berkeley: New Riders Publishing.

Kwasnik, Barbara H. 1991. The importance of factors that are not document attributes in the organisation of personal documents. *Journal of documentation* 47: 389-98.

Malone, Thomas W. 1983. How do people organize their desks? Implications for the design of office information systems. *ACM transactions on office information systems* 1n1: 99-112.

Markey, Karen. 2007a. Twenty-five years of end-user searching, Part 1: Research findings. *Journal of the American Society for Information Science and Technology* 58: 1071-81.

Mathes, Adam. 2004. Folksonomies - Cooperative classification and communication through shared metadata. *Adammathes.com*. Available <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html> (accessed May 31, 2009).

Nielsen, Marianne Lykke. 2004. Thesaurus construction: Key issues and selected readings. *Cataloging & classification quarterly* 37n3: 57-74.

Ockerbloom, John Mark. 2006. New maps of the library: Building better subject discovery tools using Library of Congress Subject Headings. *Working Paper for the CNI Task Force Meeting, December 5, 2006*. Available http://repository.upenn.edu/library_papers/48/ (accessed May 31, 2009).

Oppenheim, Charles; Morris, Anne; and McKnight, Cliff. 2000. The evaluation of WWW search engines. *Journal of Documentation* 56: 190-211.

Quintarelli, Emanuele; Resmini, Andrea; and Rosati, Luca. 2006. FaceTag: Integrating bottom-up and top-down classification in a social tagging system. *Proceedings of the 2nd European Information Architecture Conference, Berlin, Germany, September 30-October 1*. Available <http://www.facetag.org/download/facetag.pdf> (accessed May 31, 2009).

Schwartz, Candy. 2008. Thesauri and facets and tags, oh my! A look at three decades in subject analysis. *Library trends* 56 no.4: 830-42.

Shiri, Ali and Revie, Crawford. 2005. Usability and user perceptions of a thesaurus-enhanced search interface. *Journal of documentation* 61: 640-56.

Shiri, Ali and Revie, Crawford. 2006. Query expansion behavior within a thesaurus-enhanced search environment: A user-centered evaluation. *Journal of the American Society for Information Science and Technology* 57: 462-78.

Shirky, Clay. 2005. Ontology is overrated: Categories, links, and tags. *Shirky.com*. Available http://shirky.com/writings/ontology_overrated.html (accessed May 31, 2009).

Strauss, Anselm, and Corbin, Juliet. 1990. *Basics of qualitative research: Grounded theory procedures and techniques*. London: Sage.

Tang, Muh-Chyun, and Sun, Ying. 2003. Evaluation of web-based search engines using user-effort measures. *Libres* 13n2: 1-11. Available <http://libres.curtin.edu.au/libres13n2/tang.htm> (accessed May 31, 2009).

Trant, Jennifer. 2006. Exploring the potential for social tagging and folksonomy in art museums:

proof of concept. *New review of hypermedia and multimedia* 12n1: 83 - 105. Available <http://www.archimuse.com/papers/steve-nrhm-0605preprint.pdf> (accessed May 31, 2009).

Yoon, Jungwon. 2009. Towards a user-oriented thesaurus for non-domain-specific image collections. *Information processing and management* 45: 452-68.