

# Journal of Information Science

<http://jis.sagepub.com>

---

## **The folksonomy tag cloud: when is it useful?**

James Sinclair and Michael Cardew-Hall

*Journal of Information Science* 2008; 34; 15 originally published online May 31, 2007;

DOI: 10.1177/0165551506078083

The online version of this article can be found at:

<http://jis.sagepub.com/cgi/content/abstract/34/1/15>

---

Published by:

 SAGE Publications

<http://www.sagepublications.com>

On behalf of:



Chartered Institute of Library and Information Professionals

**Additional services and information for *Journal of Information Science* can be found at:**

**Email Alerts:** <http://jis.sagepub.com/cgi/alerts>

**Subscriptions:** <http://jis.sagepub.com/subscriptions>

**Reprints:** <http://www.sagepub.com/journalsReprints.nav>

**Permissions:** <http://www.sagepub.com/journalsPermissions.nav>

**Citations** (this article cites 7 articles hosted on the SAGE Journals Online and HighWire Press platforms):  
<http://jis.sagepub.com/cgi/content/refs/34/1/15>

# The folksonomy tag cloud: when is it useful?

James Sinclair and Michael Cardew-Hall

*Department of Engineering, The Australian National University, Canberra, ACT 0200, Australia*

## Abstract.

The weighted list, known popularly as a 'tag cloud', has appeared on many popular folksonomy-based websites. Flickr, Delicious, Technorati and many others have all featured a tag cloud at some point in their history. However, it is unclear whether the tag cloud is actually useful as an aid to finding information. We conducted an experiment, giving participants the option of using a tag cloud or a traditional search interface to answer various questions. We found that where the information-seeking task required specific information, participants preferred the search interface. Conversely, where the information-seeking task was more general, participants preferred the tag cloud. While the tag cloud is not without value, it is not sufficient as the sole means of navigation for a folksonomy-based dataset.

**Keywords:** collaborative tagging; folksonomies; tag clouds; user interfaces; information retrieval

## 1. Introduction

The popularity of services such as Flickr,<sup>1</sup> Delicious<sup>2</sup> and Technorati<sup>3</sup> has created a great deal of interest in tagging systems, and their potential for collaborative organization of resources. In fact, interest has been so great that articles have appeared in the popular press [1, 2], explaining the concept to audiences who would normally have little interest in information architecture. Yahoo! considered these kinds of systems important enough to make the strategic acquisition of Flickr [3] and Delicious [4]. Across the world wide web (WWW), the use of tagging systems is expanding rapidly.

### 1.1. Tagging services

Tagging services allow a participant to associate freely determined keywords (called 'tags') with a particular resource. The resource may be just about anything. Tagging services exist to tag an enormous variety of things such as photographs, URLs, podcasts, computer games, music and videos. The emergent dataset arising from all the participants' tags and resources is commonly referred to as a 'folksonomy'. Hammond et al. [5] provide an excellent review of tagging services.

---

*Correspondence to:* James Sinclair, Department of Engineering, The Australian National University, Canberra, ACT 0200, Australia. Email: james.sinclair@anu.edu.au

The concept of using freely determined keywords assigned by non-experts is not a new one [6–8]. Many scientific journals use author-generated keywords to supplement indexing, and the use of the ‘keywords’ meta-tag was an early feature of the WWW [9]. Using an uncontrolled vocabulary in this way is less costly than employing an expert to perform a domain analysis to determine an appropriate classification scheme. Allowing authors to classify their own work also avoids the ongoing cost of employing an expert to classify resources added to the repository.

There are problems, however, with allowing authors control over classification. Firstly, there is the question of expertise. The issue is summarized by Barton et al.:

The basic problem is this: who is best placed to add subject-related metadata for maximum resource discovery, the author who may know the subject area and its terminology well, or a metadata specialist, who may be better placed to step back and think about all the potential users of a resource, or about consistency of subject terms and classifications across a repository, archive or network, so that like really is classified with like. [10]

Secondly, there is the issue of deliberate misuse of metadata. There is a perception that author-generated metadata lacks credibility. Thomas and Griffin [11] argue that there is little incentive for authors to put time and effort into providing quality metadata unless it is for the purposes of self-promotion. Thus, author generated metadata is perceived as inaccurate at best, and deliberately misleading at worst. [12, 13]

Tagging services sidestep this issue of author versus expert categorization, opting instead for user created metadata [14]. That is, the people using the system are the ones who create tags, rather than system administrators or content authors. Individuals ostensibly create tags to serve their own needs, and in doing so, a consensus vocabulary emerges.

Another characteristic feature of tagging services is that they generally make a participant’s tags and resources visible to other participants. Proponents of tagging services claim that this adds a social incentive to tagging. If a participant tags a resource she is interested in, then she is helping others who share similar interests to find it too.

Proponents also suggest that this social aspect of tagging services acts as a kind of feedback mechanism for the folksonomy. When a participant observes how others have tagged a resource, they are more likely to adopt a similar tagging vocabulary when describing related resources. In this context, the folksonomy develops as an emergent vocabulary amongst participants sharing similar interests [15].

While there is much that could be said about the benefits and drawbacks of this particular approach, our purpose here is not to critique the philosophy behind collaborative tagging services. Hendry et al. use the term ‘amateur bibliography’ to refer broadly to collections of resources compiled by non-professionals for public use. They point out that proliferation of amateur bibliographies is a lasting trend, fuelled by the increasing integration of the internet into people’s lives:

A large number of bibliographies and bibliography-like information artefacts are created by non-professionals. [...] While these lists often fall short of the standards of professional bibliography, they are nevertheless important organizing instruments which, in turn, lead to a great benefit. The various forms of amateur bibliography are, in general, underappreciated and understudied. [16]

Tagging services are a fast growing subset of amateur bibliographies that warrant further investigation. A few authors have published introductory studies on tagging services [5, 17, 18], and more recently a number of papers were presented at the *Collaborative Web Tagging Workshop at WWW2006*. As stated by Marlow et al. [19], however, ‘despite the rapid expansion of applications that support tagging of resources, tagging systems are still not well studied or understood’. In this paper, we examine user interfaces for tagging services, with a particular focus on the enterprise context.

### 1.2. Tagging in the enterprise

Given the social aspects of tagging systems, there is clearly potential for these services to support knowledge sharing and collaborative work within the enterprise. The sharing of key resources and the

generation of emergent vocabulary may well help decrease the cognitive distance between people working in an organization. This helps to promote social capital, a key factor in knowledge sharing [20].

Although there is clear potential for tagging systems in the enterprise, there are still key questions to address. Some of these include:

- What factors are important to consider when scaling down a tagging system from the whole-of-internet, to small- or medium-sized organization scale [21]?
- How can we present folksonomy data in ways that enhance people's ability to locate relevant information?

In this paper, we examine aspects of these questions with particular respect to user interfaces for folksonomies.

### 1.3. Tag clouds

Tag clouds are a user interface element commonly associated with folksonomy datasets. The Wikipedia encyclopaedia contains the following entry regarding the tag cloud:

A tag cloud (more traditionally known as a weighted list in the field of visual design) is a visual depiction of content tags used on a website. Often, more frequently used tags are depicted in a larger font or otherwise emphasized, while the displayed order is generally alphabetical. Thus both finding a tag by alphabet and by popularity is possible. Selecting a single tag within a tag cloud will generally lead to a collection of items that are associated with that tag.<sup>4</sup>

In essence, the tag cloud translates the emergent vocabulary of a folksonomy into a social navigation tool [22].

Although tag clouds are extremely common among tagging services, the academic literature (at least to our knowledge) appears to be largely silent on the subject. In the blogosphere however, there has been heated discussion amongst website design and information architecture practitioners. Jeffrey Zeldman [23] initially criticized tag clouds simply for being popular. However, in the same article he described tag clouds as being 'smart' and said that this smartness is the reason 'why so many have rushed to use them'. In a subsequent article he criticized tag clouds as obscuring useful information by being skewed towards what is popular. He claims that tag clouds 'create a false intellectual equivalence between high- and low-level topics' [24]. Some of his criticisms were not so much of tag clouds, however, as folksonomies in general.

Some bloggers, such as Jon [25], wrote responses to these articles and defended the tag cloud as a navigation aid. Jon claimed that the tag cloud is a powerful tool that has yet to realize its full potential. Others, such as Porter [26], responded to the broader criticisms of folksonomies in general. Zeldman appears to claim that the popularity metric hinders access to useful items that are not tagged with popular keywords. Porter counter-argues, however, that just because an item is popular, that does not necessarily mean it is irrelevant.

Notably, Tomas Vander Wal – the man credited with coining the term 'folksonomy' [27] – described tag clouds as being 'cute but offer[ing] little value' [28].

While the discussion in the blogosphere has been lively, very few of the authors have performed empirical studies to verify their claims. As Brooks and Montanez point out:

[M]uch of the discussion regarding the advantages and disadvantages of tags and folksonomy has taken place within the blogosphere, as opposed to within peer-reviewed conferences or journals. The blogosphere has the great advantage of allowing this discussion to happen quickly and provide a voice to all interested participants, but it also presents a difficult challenge to researchers in terms of properly evaluating and acknowledging contributions that have not been externally vetted. [29]

This lack of empirical evidence prompted our investigation of the advantages and disadvantages of tag clouds for exploring folksonomy data. This fits in with our broader aim of investigating the use of tagging systems in small- to medium-sized organizations.

Table 1  
Participants

Participants:	89
Female:	6
Male:	83
Age range:	18–24
Participants with a blog:	10
Participants who use a social book-marking service (e.g. Delicious, Furl, Magnolia):	2
Articles tagged:	973
Mean articles tagged per session:	108.1
Mean articles tagged per participant:	10.9

#### 1.4. The study

The question we addressed here is ‘Do tag clouds provide value to people seeking information from a folksonomy dataset?’ To answer this question we created a folksonomy-like dataset, and asked participants to perform information-seeking tasks on it. We provided participants with the option of using either a tag cloud or a search box interface to query the database. Our hypothesis was that if people found no value in a tag cloud for information seeking, then they would primarily use the search box to seek information. At best, people would attempt one or two queries using the tag cloud, then resort to the search box.

We found that while some participants preferred to use the search box exclusively, a significant proportion of participants used the tag cloud to find information. In total, the number of questions answered using the tag cloud was greater than those answered using the search box. Thus, we believe that the tag cloud does provide value for navigating a folksonomy dataset.

Examining survey comments and usage patterns of the data, we explore the benefits and limitations of the tag cloud.

## 2. Method

A key characteristic of the folksonomy is that it is *user created metadata* [14], that is, the people who use the system are also the ones tagging items. Hence, instead of downloading an existing folksonomy dataset for use in this study [as with other studies such as 15, 18, 30], we asked participants in the study to tag articles from Slashdot Science<sup>5</sup> to create a folksonomy-like dataset. This ensured that the participants querying the dataset had also contributed tags to the dataset.

To conduct the experiment, we recruited participants from students enrolled in a first-year engineering class offered by the Department of Engineering at the Australian National University. Table 1 summarizes other characteristics of the participants.

In the experiment, we asked participants to tag around 10 articles each, using the interface shown in Figure 1. The system selected articles randomly from a pool of 1074 articles, dated from 17 August 2005 to 25 May 2006. Explaining the experiment and allowing participants to tag 10 articles took approximately 25 minutes for each session.

The system allowed participants to enter as many tags as they wished, separated by commas. This made it possible to use spaces in tags, rather than restricting the participant to a single word. This is in contrast to many popular tagging systems such as Delicious and Flickr, which only allow single word tags. We allowed the use of multi-word tags to eliminate the problem of establishing a convention for word combination. One study found that at least 10% of tags in Delicious were compound words created ‘by putting a symbol or piece of punctuation inside the tag to represent a space’, yet no consensus or convention had been chosen by the Delicious user community [17]. Allowing spaces (in a manner similar to that adopted by RawSugar<sup>6</sup>) avoids this issue.

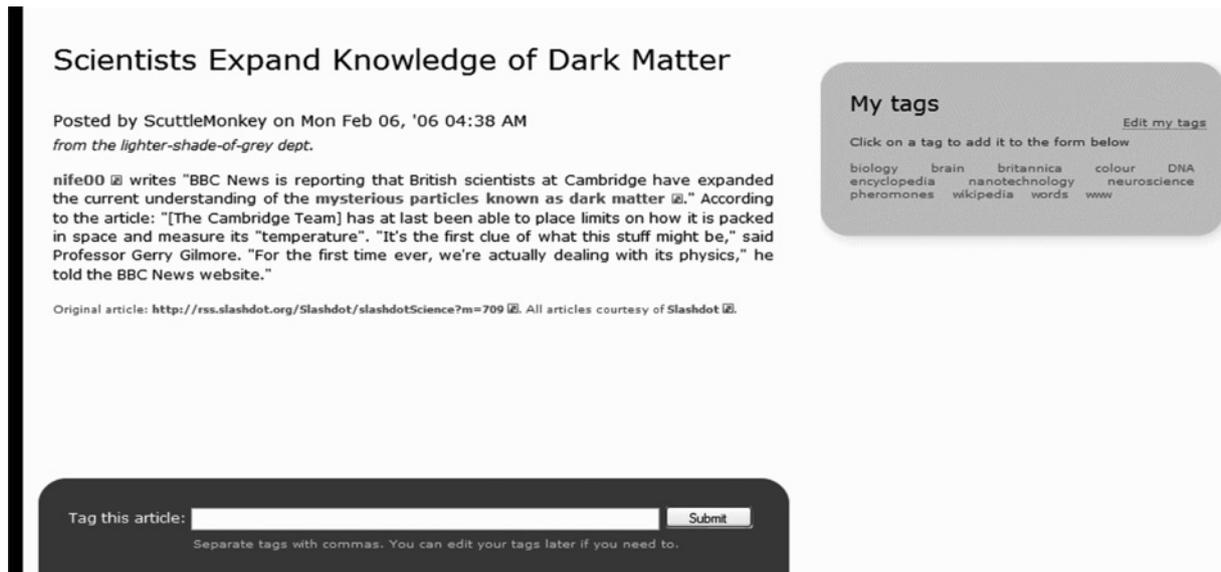


Fig. 1. The interface for tagging articles.

The tagging interface also listed all the tags a participant had used to date on the right-hand side of the screen under the heading 'My tags'. Clicking on one of the tags added the tag to the form at the bottom of the screen. This is a timesaving feature common in tagging services such as Delicious and CiteULike.<sup>7</sup> We included this feature to encourage consistency in tagging by reminding participants of tags they had used previously. For example, a participant might enter 'web design' for one article and 'website design' for another. We included the 'My tags' to limit this individual non-conformity.

Each session was divided into two halves, with the first half being devoted to tagging articles, and the second half devoted to querying the database of tagged articles. In the second half, we asked participants to answer 10 questions (see Table 2) by querying the collection of tagged articles. The questions were designed to elicit a range of information-seeking behaviour, ranging from general browsing to specific information seeking. Questions 2 and 4 were the most general, giving no specific topic to search on, while questions 1 and 6 related to specific topics represented by a relatively small number of articles in the dataset.

To query the database, we presented participants with the option of using either a tag cloud, or a conventional search box. Both options were presented on the same screen, with approximately half of the horizontal screen space being given to each (see Figure 2). The system recorded which interface participants used to make their queries, and which question was displayed on the screen at the time.

The tag cloud displayed the top 70 most-used tags in the database. As folksonomy frequency tends to follow a power law distribution [15], we made the size of each tag proportional to the log of its frequency. The specific equation used was:

$$\text{TagSize} = 1 + C \frac{\log(f_i - f_{\min} + 1)}{\log(f_{\max} - f_{\min} + 1)} \quad (1)$$

where  $f_i$  is the frequency of tag  $i$ ,  $f_{\min}$  is the minimum frequency in the top 70 tags,  $f_{\max}$  is the maximum frequency in the top 70 tags, and  $C$  is a scaling constant that determines the maximum text size. This gives a tag size relative to the default font size set by the browser.

Clicking on a tag returned a list of all articles tagged with that keyword. The results page for the tag cloud interface showed a subsequent tag cloud for all the articles featuring the selected tag (Figure 3). Clicking on a tag in this sub-cloud would further refine the results to articles containing both tags, and so on for sub-sub-clouds, etc.

Table 2  
Questions asked of participants to elicit information-seeking behaviour

1. Tell me about an interesting application of nanotechnology.
2. Paste the title of an article you find interesting.
3. Paste the title of an article that contains something about Mars.
4. Paste the title of an article you find amusing.
5. Paste the title of an article related to biology.
6. Tell me something about sustainable fuels.
7. Tell me about a scientific idea/discovery that has military applications.
8. Paste the title of an article related to NASA.
9. Paste the title of an article related to archaeology.
10. Paste the title of an article related to video game addiction.

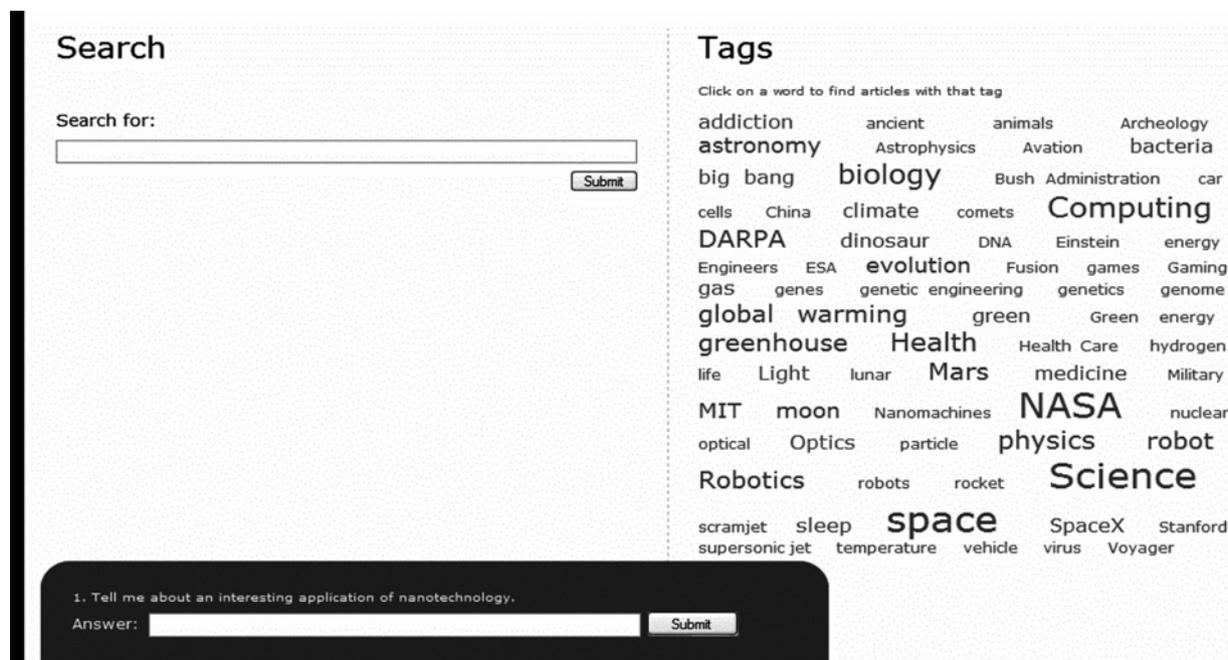


Fig. 2. The search interface showing the search box and tag cloud. This particular screenshot is from the third session.

The system generated query results for the search box using the Ranking with Vector Spaces algorithm provided by MySQL,<sup>8</sup> using the full text of each article. This algorithm uses a variation on the well known TF-IDF term-weighting method. In this method the database system assigns each term in the collection a weight, according to its frequency in the collection and its frequency in each article. When a participant performs a query, the algorithm selects and ranks articles according to the matched query terms and their individual weights.

By contrast, query results from the tag cloud were generated by simply retrieving all articles tagged with the chosen keyword. That is, any article tagged with the relevant keyword was included in the search results. This left the question of how to rank the search results, however. In order to maintain consistency, we ranked the result set displayed by the tag cloud using the same Ranking with Vector Spaces algorithm used for the search box. So, while the methods used to *select* which articles to display were different for the search box and tag cloud, the same method was used to *rank* the displayed results.

While most folksonomy-based websites order results by how recently items were tagged, and the number of people tagging the item, in this system such an organization would have had little relevance. Other methods of ranking folksonomy search results have been proposed, e.g. [31]; however,

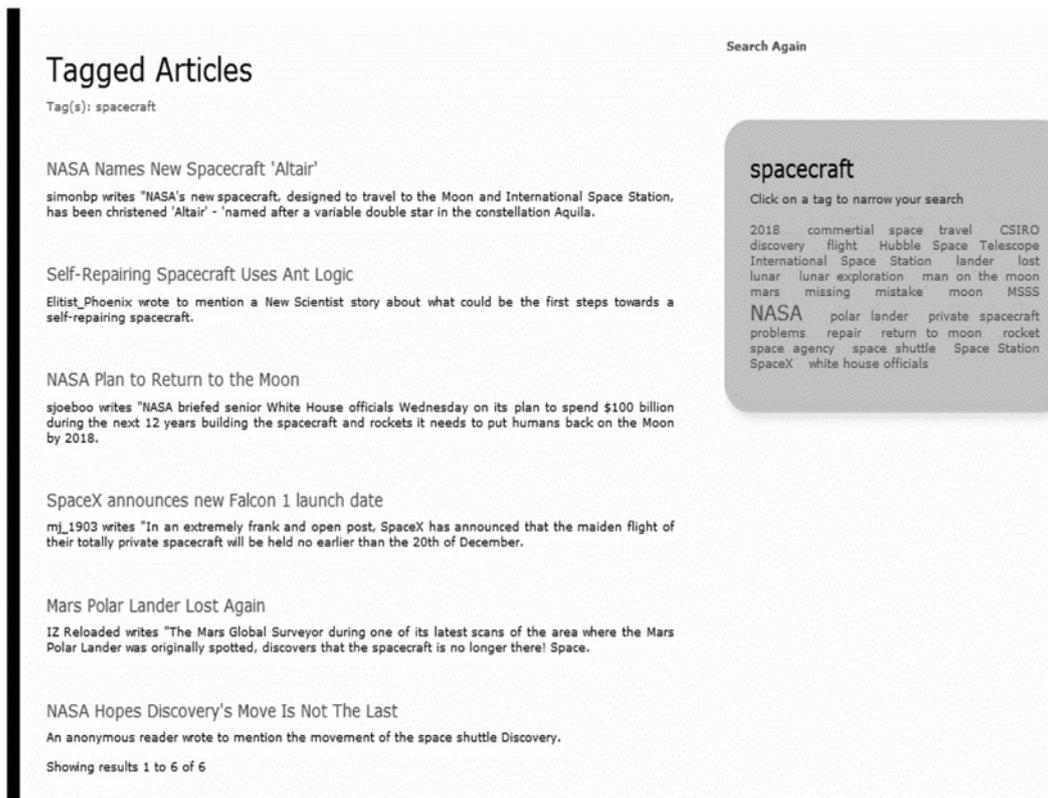


Fig. 3. Example of results from a tag cloud query.

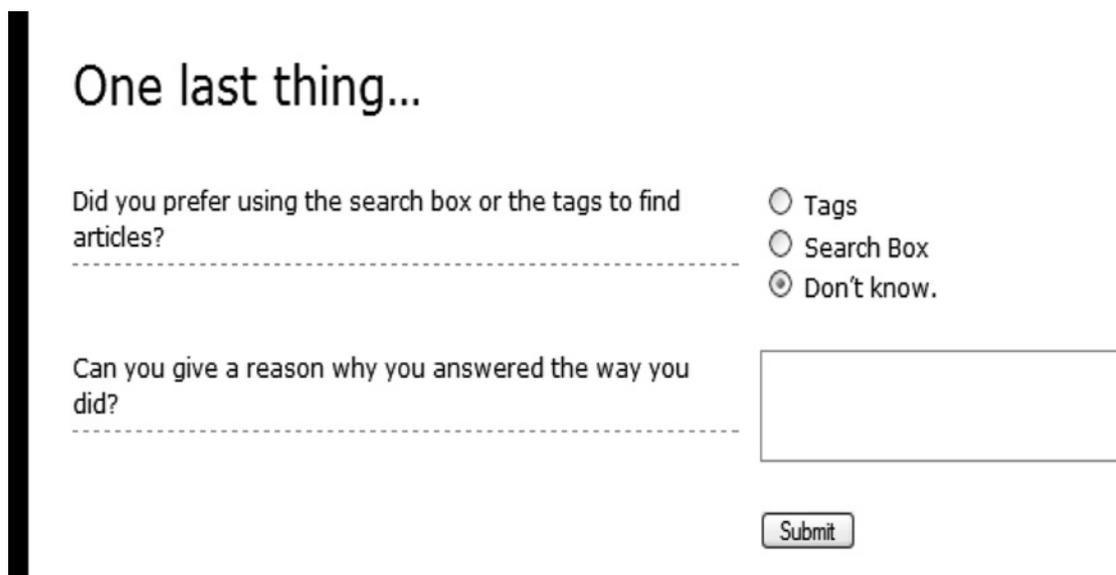


Fig. 4. The exit survey.

our intention was not to compare the precision and recall of the two techniques, but rather to examine user perceptions and usage patterns. Hence, we attempted to maintain consistency by ranking results using the same algorithm for both query methods.

After completing 10 questions, the system presented participants with an exit survey asking two questions (see Figure 4).

### 2.1. *Characteristics of the dataset*

The folksonomy dataset created through this process is quite different to that of frequently studied, large-scale systems such as Flickr or Delicious. The experimental conditions here, however, reflect contexts that often occur in an enterprise context. Not all organizations have thousands of employees. Managers often delegate data entry and information management tasks to subordinates. There is potential in many organizational settings for people to be tagging resources that they did not select, without any personal benefit to themselves. Of course, such situations would not be ideally suited to a folksonomy; in practice, however, all methods of categorization are less than ideal [32, 33].

The specific differences, in contrast to a larger-scale folksonomy, are as follows:

- The number of participants was relatively small, when compared with studies of tagging systems in the enterprise [such as 21, 34, 35] which report upwards of 300 participants.
- Participants did not select articles to tag. This is in contrast to most tagging services where participants select for themselves which resources they wish to tag. This provides a kind of quality filter not present in the experiment we conducted.
- As the 'shared' dataset was artificially selected, participants lacked a context for tagging. Participants in most tagging systems find some kind of personal benefit for tagging. Delicious, for example, allows participants to better organize their own bookmarks. Thus, tags created in this context have personal relevance for the participant. In this case, participants were tagging simply because it was part of the experiment.
- There was no social feedback aspect to tagging. Unlike most tagging systems, participants were not able to browse others' tags in this experiment. Since participants did not select which articles to tag, this would have had little relevance. However, this prevented participants from observing how others tagged articles, potentially resulting in less convergence on tag usage.

These conditions produce something of a worst-case-scenario folksonomy. However, if a tag cloud shows value for navigation in a less-than-ideal dataset such as this, then we assume that it will perform even better with more typical folksonomy data.

## 3. Results

### 3.1. *Last strike queries*

Do tag clouds provide value to people seeking information from a folksonomy dataset? To examine this question, we assume that the last query made before answering a question provided the relevant information to answer. We shall refer to the last query made before answering as a 'last-strike query'.

In practice, this assumption about last-strike queries may not always hold true. Someone may continue querying a database, even after they have found a relevant resource, simply to see if a 'better' resource is available. In the context of this study however, the majority of questions asked participants to 'copy and paste' from articles. This makes it more likely that the last-strike query did indeed provide the relevant answer.

Each of the 89 participants answered 10 questions, giving 890 last-strike queries. We assume that those who did not search to answer a question either answered from their own knowledge, or used the results from the previous search.

Table 3 shows that participants answered the majority of questions using the tag cloud. Thus, it seems that the tag cloud does provide some kind of value. Examining the last strike queries for each question, however (see Figure 5), participants favoured the search box for six of the 10 questions. Investigating this further, we find two scenarios where use of the tag cloud outweighed use of the

Table 3  
Last-strike queries

Method:	Tag cloud	Search box	Did not search
Last-strike queries:	427 (48.0%)	367 (41.2%)	96 (10.8%)

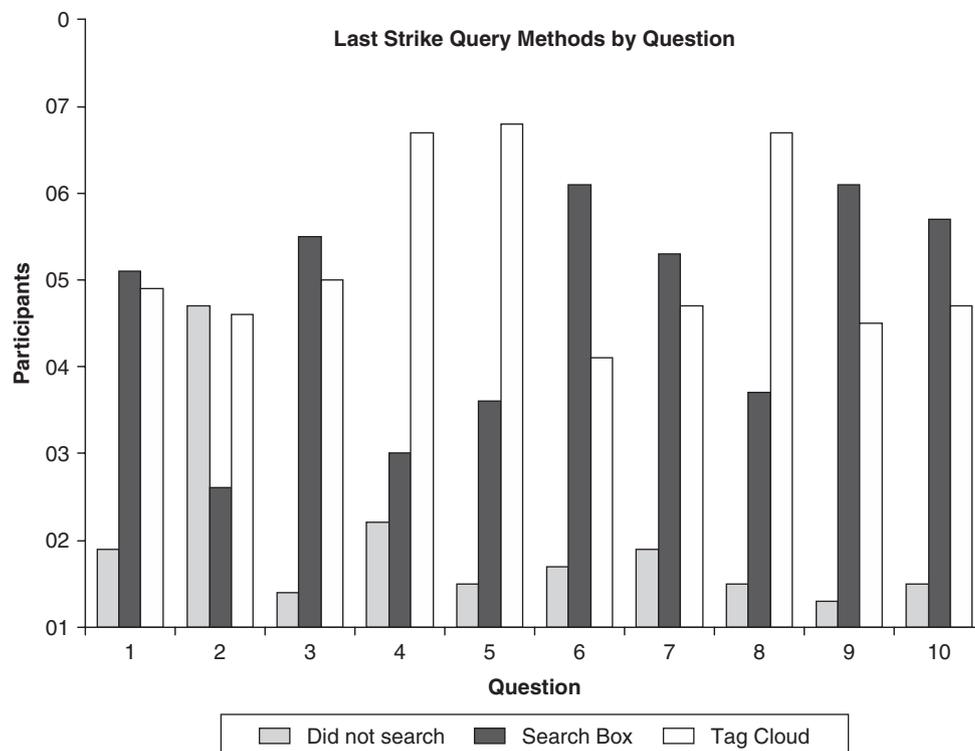


Fig. 5. Query method used to answer each question.

search box. The first was where the information-seeking task was broad and non-specific, as in questions 2 and 4. The second was when the tag cloud contained a keyword relevant to the question.

### 3.2. Browsing versus finding

Questions 2 and 4 were broad, non-specific questions:

2. Paste the title of an article you find interesting.
4. Paste the title of an article you find amusing.

The number of last strike queries for the tag cloud was much greater than the search box for both these questions. This suggests that tag clouds are useful when the information-seeking task is non-specific. That is, tag clouds support browsing or serendipitous discovery.

Other authors have observed that this is a feature of folksonomies in general. Mathes [14] commented that one of the strengths of folksonomies is serendipity: 'Browsing the system and its inter-linked related tag sets is wonderful for finding things unexpectedly in a general area.' Brooks and Montanez [29] also found that tagging is most effective at placing articles into broad categories, but is less effective as a tool for indicating an article's specific content.

Table 4  
Mean queries to answer question where participants relied on a single interface

	No keyword in cloud	Keyword in cloud
Search Box	1.53	1.18
Tag Cloud	2.05	1.31

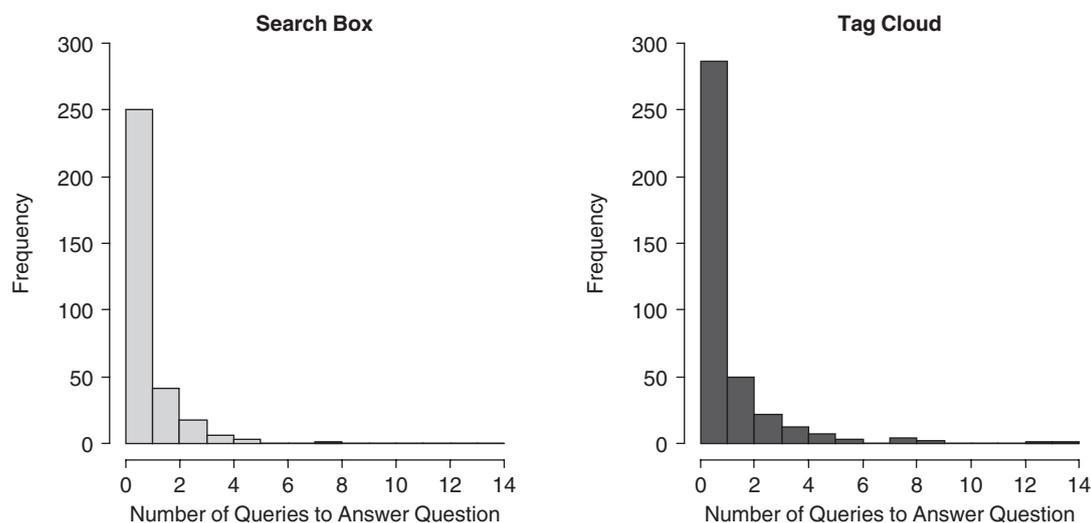


Fig. 6. Number of queries required to answer questions where participants relied solely on a single interface.

In addition, participants' comments from the exit survey also support this observation. Eight respondents indicated that their choice of interface depended on the information-seeking task, often commenting that the tags supported more general browsing. Ten participants commented that the search box allowed for greater specificity when entering queries. Examples included:

[...] the tags were more usefull [sic] for finding things about general topics. [...] the search bar was better for spicific [sic] things

[...] the search box lets you find exactly what you want, however if you wanted to find related articles, the tags are more useful

If searching for something specific you can specify more constraints about the subject so less results are returned, otherwise for more general use the tags would be good

Conversely, it seems that the tag cloud was less useful when the information-seeking task was more specific.

In general, it appears that answering a question using the tag cloud required more queries per question than the search box. We cannot directly compare the mean number of searches for each interface, however, since a participant may have used a mixture of both methods. To compare number of searches to find an answer, we examine the cases where a participant solely relied on just one interface to answer a question.

The histograms for both interfaces (Figure 6) show that most questions were answered with a single query. In general however, participants who relied solely on the tag cloud made more queries when answering a question, even when a relevant keyword was present in the tag cloud (see Table 4). The 'tail' of the tag cloud histogram in Figure 6 shows two cases where participants made 13 to 14 queries before answering a question. This suggests that participants who used the tag cloud were more likely to engage in browsing behaviour.

Table 5  
Last strike queries when relevant keywords were present in tag cloud

Keyword in tag cloud:	No		Yes	
Did not search:	65	(14.3%)	31	(7.1%)
Search Box:	194	(42.8%)	173	(39.6%)
Tag Cloud:	194	(42.8%)	233	(53.3%)

Table 6  
Participants' preferences from the exit survey

Search	Don't know	Tag	Did not answer
35 (39.33 %)	11 (12.36 %)	26 (29.21 %)	17 (19.10 %)

### 3.3. Presence of relevant keywords in the tag cloud

The other questions where use of the tag cloud dominated were Questions 5 and 8. These were more specific questions than 2 and 4, referring to a particular subject or topic:

5. Paste the title of an article related to biology.
8. Paste the title of an article related to NASA.

With these questions, however, there was a relevant keyword present in the tag cloud for the majority of sessions. The keyword *NASA* was present in the tag cloud for every session. The keyword *biology* was present in the tag cloud for seven of the nine sessions. When present in the tag cloud, both keywords were relatively large in proportion to the surrounding text. A preference for the tag cloud with these questions makes sense, since clicking on a single keyword requires less effort than typing search terms and submitting a form.

Investigating this further, we find that last-strike usage of the tag cloud increased whenever there was a relevant keyword present in the tag cloud (see Table 5).

This is consistent with participant comments that clicking on a single keyword was faster than typing it into the search box. For example:

It was quicker, instead of typing in a search word/s, you can just click on the provided keywords to find a relevant article.

[...] clicking is fast[er] than typing [...]

When the question had a tag associated with it, [it] was faster to use the tag

### 3.4. Participants' preference

In the exit survey, a majority of participants stated a preference for the search box (Table 6). Participants' reasons for their preference were coded and tabulated as shown in Table 7. There was a wide variety of responses; however, the top three reasons are particularly interesting since they reflect our findings regarding 'browsing versus finding'.

A factor to consider here is the makeup of the participant group. As mentioned previously, we recruited participants from a first-year engineering class. Amongst such a group there is likely to be a high level of computer literacy and familiarity with information seeking tools such as search engines. At the same time, very few of the participants appeared to be familiar with the concept of tagging or tag clouds. Only two participants indicated that they used a social bookmarking service,

Table 7  
Reasons for choosing one interface over another

Responses	Reason category
10	Search box allowed greater specificity
8	Depends on information-seeking task
7	The tag cloud was faster
4	The search box was faster
4	The tag cloud 'suggests' useful query terms
4	The search box was easier to use (or the tag cloud was more difficult)
4	The search box is familiar or the tag cloud unfamiliar
3	The tag cloud was easier to use (or the search box was more difficult)

and only 10 of the 89 participants indicated that they kept a blog. This may have led to a greater preference for the search box due to its familiarity. Comments from the exit survey support this. For example:

I am used to searching with search boxes on databases and the internet

[I preferred the search box because I am] used to it

[The search box is] a system I am used to

### 3.5. Tag cloud as a visual summary

There is evidence to suggest that formulating a free text search query requires greater mental activity than browsing a list of links [36]. In order to formulate a query, a participant must *recall* relevant words to enter into the search box – a cognitively more expensive task than *recognizing* relevant words in a list [37]. A number of the comments from the exit survey indicated that the tag cloud 'suggested' useful search terms, saving the participant some cognitive effort. Examples of participant's comments include:

[T]he tags made it so you didn[']t need to think of what you wanted to find.

The tags invoked ideas immediately in my mind, making access easier

When you want to search something but you don't know where to start, using tags are the right way to start.

In the experiment sessions, we also observed that participants for whom English was a second language seemed to find the tag cloud particularly helpful in this regard. This suggests that tag clouds may have the potential to increase accessibility for people performing information-seeking tasks in a second language. More research is needed to confirm this hypothesis, however.

### 3.6. Tag cloud occlusion

One disadvantage of the tag cloud is that when used in this manner, it does not allow direct access to all articles in the database. When the system presented search results from a tag based query, it also showed a cloud of co-occurring tags as displayed in Figure 3. Clicking on one of these co-occurring tags presented articles tagged with *both* keywords. Thus, if an article had no tags in the front tag cloud (that is, in the 70 most popular tags), then it was inaccessible from the tag cloud interface.

In every session of the experiment, a significant number of articles were not accessible from the tag cloud. Interestingly, while the number of tags in the cloud remained constant (70 tags), the number of occluded articles increased proportionally to the total number of articles (see Figure 7). The percentage of occlusion remained roughly constant at about 55.5% for each session (see Figure 8). This inability to make all articles accessible from a single page is a major disadvantage compared with the search box. Hence, tag clouds are clearly not sufficient as the sole means of navigating a folksonomy dataset.

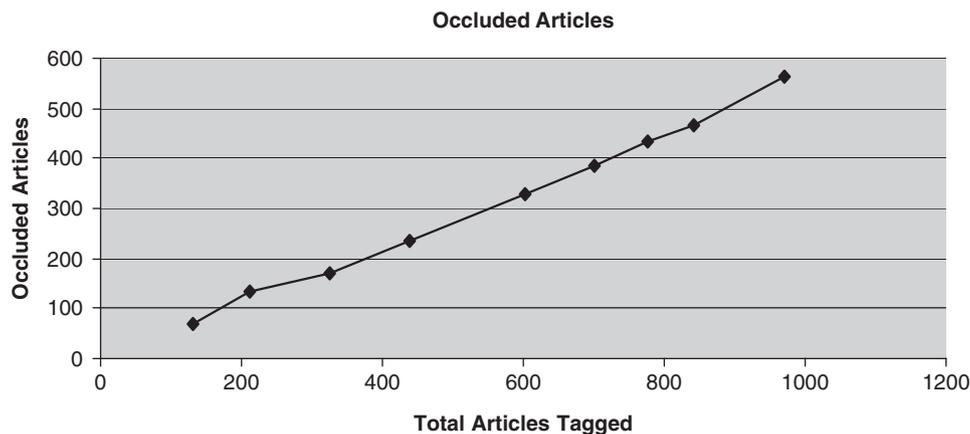


Fig. 7. Articles inaccessible from the tag cloud.

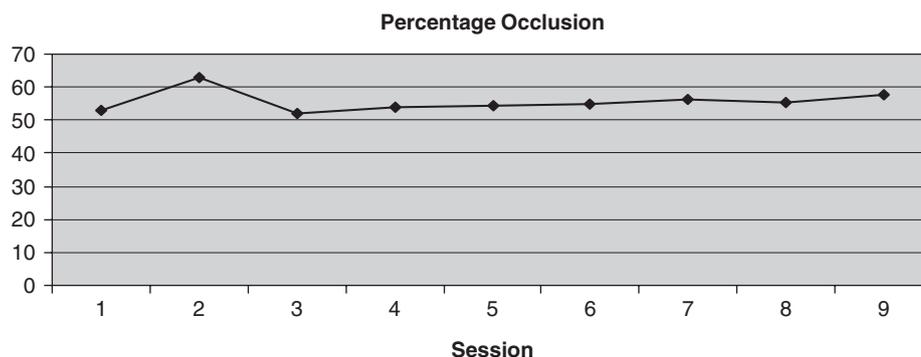


Fig. 8. Percentage of articles not accessible from the tag cloud.

## 4. Conclusion

This paper reports on an attempt to answer the question: do tag clouds provide value to people seeking information from a folksonomy dataset? After observing that a majority of questions were answered using the tag cloud, we conclude that the tag cloud does indeed provide value. The tag cloud has a number of positive attributes:

- It is particularly useful for browsing or non-specific information discovery. This is partly because tags are best suited to broad categorization, as opposed to specific characterization of content. The tag cloud also reduces the cost of a query, since a single click requires less effort than typing search terms and submitting a form.
- The tag cloud provides a visual summary of the contents of the database. A number of participants commented that this gave them an idea of where to begin their information seeking. The tag cloud provides potential for a two-phase search process where: (1) browsing activity using the tag cloud familiarizes the information seeker with the domain; and (2) a search mechanism allows more focussed queries based on the information gained from browsing.
- It appears that scanning the tag cloud requires less cognitive load than formulating specific query terms. That is, scanning and clicking is 'easier' than thinking about what query terms will produce the best result and typing them into the search box. We hypothesize that this may be particularly useful for people seeking information in a second language.

The tag cloud is not without its limitations however:

- A corollary of the tag cloud's suitability for general browsing is its unsuitability for seeking specific information. Many of the participants commented that the tag cloud did not allow them to narrow their search enough to answer the given questions.
- For this experiment, we found that roughly half the articles in the dataset were inaccessible from the tag cloud. Most tagging systems mitigate this by including tag links when viewing individual articles, thus exposing some of the less popular tags. However, this is not necessarily useful when someone is seeking specific information.

In light of these disadvantages, it is clear that the tag cloud is not sufficient as the sole means of navigating a folksonomy dataset. Some form of search facility or supplementary classification scheme is necessary to expose all the articles. On the other hand, it would be unfair to say that the tag cloud offers no value as a navigation aid. Where relevant keywords are present, clicking on a tag is quicker and easier than entering query terms into a search box. It also provides a useful visual summary of the contents of the database.

## Endnotes

- 1 <http://flickr.com/>
- 2 <http://del.icio.us/>
- 3 [www.technorati.com/](http://www.technorati.com/)
- 4 [http://en.wikipedia.org/w/index.php?title=Tag\\_cloud&oldid=63880649](http://en.wikipedia.org/w/index.php?title=Tag_cloud&oldid=63880649) (accessed 17 July 2006).
- 5 <http://slashdot.org/science>
- 6 [www.rawsugar.com](http://www.rawsugar.com)
- 7 [www.citeulike.org](http://www.citeulike.org)
- 8 See <http://dev.mysql.com/doc/internals/en/full-text-search.html> for details.

## References

- [1] O. Burkeman, Folksonomy, *The Guardian* 12 September (2005).
- [2] D.H. Pink, Folksonomy, *The New York Times* 11 December (2005).
- [3] G. Price, *Yahoo Acquires Flickr* (2005). Available at: <http://blog.searchenginewatch.com/blog/050320-213044> (accessed 5 June 2006).
- [4] M. Arrington, *Yahoo.icio.us? – Yahoo Acquires Del.icio.us* (2005). Available at: [www.techcrunch.com/2005/12/09/yahoo-acquires-delicious/](http://www.techcrunch.com/2005/12/09/yahoo-acquires-delicious/) (accessed 5 June 2006).
- [5] T. Hammond, T. Hannay, B. Lund and J. Scott, Social bookmarking tools (I): a general review, *D-Lib Magazine* 11(4) (2005). Available at: <http://dx.doi.org/10.1045/april2005-hammond> (accessed 7 October 2006).
- [6] E. Svenonius, Unanswered questions in the design of controlled vocabularies, *Journal of the American Society for Information Science* 37(5) (1986) 331–40.
- [7] E. Tonkin, Folksonomies: the fall and rise of plain-text tagging, *Ariadne* 47 (2006). Available at: [www.ariadne.ac.uk/issue47/tonkin/](http://www.ariadne.ac.uk/issue47/tonkin/) (accessed 19 July 2006).
- [8] J. Greenberg, M.C. Pattuelli, B. Parsia and W.D. Roberts, Author-generated Dublin Core metadata for web resources: a baseline study in an organization, *Journal of Digital Information* 2(2) (2001). Available at: <http://jodi.ecs.soton.ac.uk/Articles/v02/i02/Greenberg> (accessed 5 June 2006).
- [9] P. Miller, Metadata for the masses, *Ariadne* 5 (1996). Available at: [www.ariadne.ac.uk/issue5metadata-masses](http://www.ariadne.ac.uk/issue5metadata-masses) (accessed 5 June 2006).
- [10] J. Barton, S. Currier and J.M.N. Hey, Building quality assurance into metadata creation: an analysis based on the learning objects and e-prints communities of practice. In: S. Sutton et al. (eds), *DC-2003: Proceedings of the International DCMI Metadata Conference and Workshop, Seattle, Washington, USA, 2003* (Information Institute of Syracuse, Syracuse, NY, 2003).

- [11] C.F. Thomas and L.S. Griffin, Who will create metadata for the Internet?, *First Monday* 3(12) (1999). Available at: [www.firstmonday.dk/issues/issue3/\\_12/thomas/index.html](http://www.firstmonday.dk/issues/issue3/_12/thomas/index.html) (accessed 26 October 2006).
- [12] C. Doctorrow, *Metacrap: Putting the Torch to Seven Straw-Men of the Meta-Utopia* (2001). Available at: [www.well.com/~doctorow/metacrap.htm](http://www.well.com/~doctorow/metacrap.htm) (accessed 27 October 2006).
- [13] E. Fidelman, *Metadata quality and the use of hierarchical schemes to determine meta keywords: an exploration* (Master's dissertation, University of North Carolina, Chapel Hill, 2006).
- [14] A. Mathes, *Folksonomies – Cooperative Classification and Communication Through Shared Metadata* (2004). Available at: <http://adammathes.com/academic/computer-mediated-communication/folksonomies.html> (accessed 10 March 2006).
- [15] C. Cattuto, V. Loreto and L. Pietronero, *Collaborative Tagging and Semiotic Dynamics* (2006). Arxiv preprint cs.CY/0605015.
- [16] D.G. Hendry, J.R. Jenkins and J.F. McCarthy, Collaborative bibliography, *Information Processing & Management* 42(3) (2006) 805–25.
- [17] M. Guy and E. Tonkin, Tidying up tags? *D-Lib Magazine* 12(1) (2006). Available at: <http://dx.doi.org/10.1045/january2006-guy> (accessed 3 June 2006).
- [18] S.A. Golder and B.A. Huberman, Usage patterns of collaborative tagging systems, *Journal of Information Science* 32(2) (2006) 198–208.
- [19] C. Marlow, M. Naaman, D. Boyd and M. Davis, HT06, tagging paper, taxonomy, Flickr, academic article, to read. In: U.K. Wil et al. (eds), *Proceedings of the 17th Conference on Hypertext and Hypermedia*, (ACM, New York, 2006) 31–39.
- [20] M. Huysman and V. Wulf, IT to support knowledge sharing in communities: towards a social capital analysis, *Journal of Information Technology* 21(1) (2006) 40–51.
- [21] D. Millen, J. Feinberg and B. Kerr, Social bookmarking in the enterprise, *Queue* 3(9) (2005) 28–35.
- [22] A. Dieberger, P. Dourish, K. Höök, P. Resnick and A. Wexelblat, Social navigation: techniques for building more usable systems, *Interactions* 7(6) (2000) 36–45.
- [23] J. Zeldman, *Tag Clouds are the New Mullets* (2005). Available at: [www.zeldman.com/daily/0405d.shtml](http://www.zeldman.com/daily/0405d.shtml) (accessed 24 June 2006).
- [24] J. Zeldman, *Remove Forebrain and Serve* (2005). Available at: [www.zeldman.com/daily/0505a.shtml](http://www.zeldman.com/daily/0505a.shtml) (accessed 23 June 2006).
- [25] N. Jon, *Tag Clouds: a Response* (2005). Available at: [www.nicholasjon.com/archives/2005/5/4/tag\\_clouds\\_a\\_response](http://www.nicholasjon.com/archives/2005/5/4/tag_clouds_a_response) (accessed 18 July 2006).
- [26] J. Porter, *Read This Post Because It Has Zeldman in the Title* (2005). Available at: <http://bokardo.com/archives/zeldman-in-the-title> (accessed 25 July 2006).
- [27] T. Vander Wal, *Folksonomy Definition and Wikipedia* (2005). Available at: [www.vanderwal.net/random/entrysel.php?blog=1750](http://www.vanderwal.net/random/entrysel.php?blog=1750) (accessed 25 July 2006).
- [28] T. Vander Wal, *Memephoria* (2006). Available at: [www.vanderwal.net/random/entrysel.php?blog=1829](http://www.vanderwal.net/random/entrysel.php?blog=1829) (accessed 27 June 2006).
- [29] C.H. Brooks and N. Montanez, Improved annotation of the blogosphere via autotagging and hierarchical clustering. In: L. Carr et al. (eds), *WWW2006: Proceedings of the 15th International Conference on World Wide Web, Edinburgh, Scotland*, (ACM Press, New York, 2006) 625–31.
- [30] R. Lambiotte and M. Ausloos, *Collaborative Tagging as a Tripartite Network* (2005). Arxiv preprint cs.DS/0512090.
- [31] A. Hotho, R. Jäschke, C. Schmitz and G. Stumme, Information retrieval in folksonomies: search and ranking. In: Y. Sure and J. Domingue (eds), *The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, 2006* (Springer, Heidelberg, 2006) 411–42.
- [32] G.C. Bowker and S.L. Star, *Sorting Things Out: Classification and Its Consequences (Inside Technology)*, (MIT Press, Cambridge, MA, 1999).
- [33] G. Lakoff, *Women, Fire, and Dangerous Things* (University Of Chicago Press, Chicago, 1990).
- [34] L. Damianos, J. Griffith, D. Cuomo, D. Hirst and J. Smallwood, Onomi: social bookmarking on a corporate intranet. In: *Collaborative Web Tagging Workshop at WWW2006, Edinburgh, Scotland, 2006*. Available at: [www.topixa.com/www2006/28.pdf](http://www.topixa.com/www2006/28.pdf) (accessed 3 July 2006).
- [35] S. Farrell and T. Lau, Fringe contacts: people-tagging for the enterprise. In: *Collaborative Web Tagging Workshop at WWW2006, Edinburgh, Scotland, 2006*. Available at: [www.rawsugar.com/www2006/25.pdf](http://www.rawsugar.com/www2006/25.pdf) (accessed 3 July 2006).
- [36] G. Marchionini, *Information Seeking in Electronic Environments* (Cambridge University Press, Cambridge, 1995).
- [37] P.A. Nobel and R.M. Shiffrin, Retrieval processes in recognition and cued recall, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27(2) (2001) 384–413.