

Freetagging

Paul Hoffman <hoffman@simmons.edu>
LIS 500 – Summer 2005

Introduction

This paper will explore the increasingly common phenomenon of *social bookmarking*, also known as *freetagging*. It will examine what freetagging is, where it came from, the uses to which it's being put, and some of the details of how it works. It will attempt to place the use of tags – the uncontrolled, user-assigned keywords that users of freetagging services create to categorize or otherwise label Web resources – in the larger context of metadata systems and the Web by comparing freetagging with conventional metadata practices and the Semantic Web. Finally, the paper will conclude by exploring trends and possible future directions for freetagging, with an emphasis on the role that library professionals might play in the future of this exciting – and sometimes daunting – new phenomenon.

What is freetagging?

Recently, Web sites such as del.icio.us¹ and Furl² have captured the attention and imagination of Internet architects, bloggers, Internet-related news sites, and even major newspapers. (See, for example, Biddulph, 2004; Green and Hof, 2005; and Markoff, 2005.) These Web sites – along with similar sites such as the photo-sharing service Flickr³ and a number of others that have sprung up in their wake – provide free bookmarking services that enable users to save images or links to Web pages that interest them, much as they might save a bookmark to a Web page in their favorite browser. What distinguishes them from the bookmarking capabilities built into today's Web browsers, and from older online bookmarking services that let users organize their bookmarks in hierarchical folders, is that the freetagging sites let users associate any number of *tags* – keywords of their own choosing – with each saved image or bookmarked Web page. The resulting bookmark or *tag record* – containing the image or link, the user's ID, and the tags they have provided – is stored on the server for anyone to see.⁴

To use a freetagging service, one begins by registering on its Web site; this is a familiar process that entails filling out a simple form with a unique user name (an ID or 'handle'), a password, and, in some cases, an e-mail address.⁵ After registering and logging in for the first time, the user may begin tagging Web resources using the service.⁶

To encourage users to tag Web pages, nearly all freetagging services provide one or more bookmarklets – snippets of Javascript embedded in a hyperlink – that make the tagging process quick and easy. Saving the bookmarklet in one's browser is generally as simple as dragging the

¹ <http://del.icio.us/>

² <http://www.furl.net/>

³ <http://flickr.com/>

⁴ Many freetagging services allow users to create private tag records that only they can see, but this doesn't appear to be a common practice. Some also save a copy of the pages the user has tagged, to ensure that the page's content is available for future retrieval.

⁵ The e-mail address is typically used only to recover or reset passwords.

⁶ Most freetagging sites use browser cookies to identify users, eliminating the need for repeated logins.

link that contains it from a page in the freetagging site to the browser's bookmarks bar or to a folder in the user's list of saved bookmarks. Because of differences among browsers, these Javascript links are sometimes browser-specific, so a given freetagging site may provide a number of different bookmarklets.

The tagging process itself is quick and easy. While viewing a Web page, the user clicks on the appropriate bookmarklet; the Javascript code within it opens a window containing a tag creation form.⁷ This form includes fields for the resource's URL and title – these fields are already filled in by the bookmarklet but may be edited – and fields for tags and free-text comments.⁸ The user completes the form and submits it to the freetagging site, which creates a tag record containing the user's ID and the contents of the form. The tagging window closes and the user resumes what she was doing. The new tag record may be retrieved at any time by going back to the freetagging site, as described below.

Figure 1 shows the moment when the user has decided to tag a Web page and is clicking on the 'Tag it' bookmarklet that begins the process. **Error! Reference source not found.** shows the resulting window containing the fields to be filled out; the user has just entered the three tags *cat*, *dog*, and *funny* and is about to click on the 'Tag It' button.

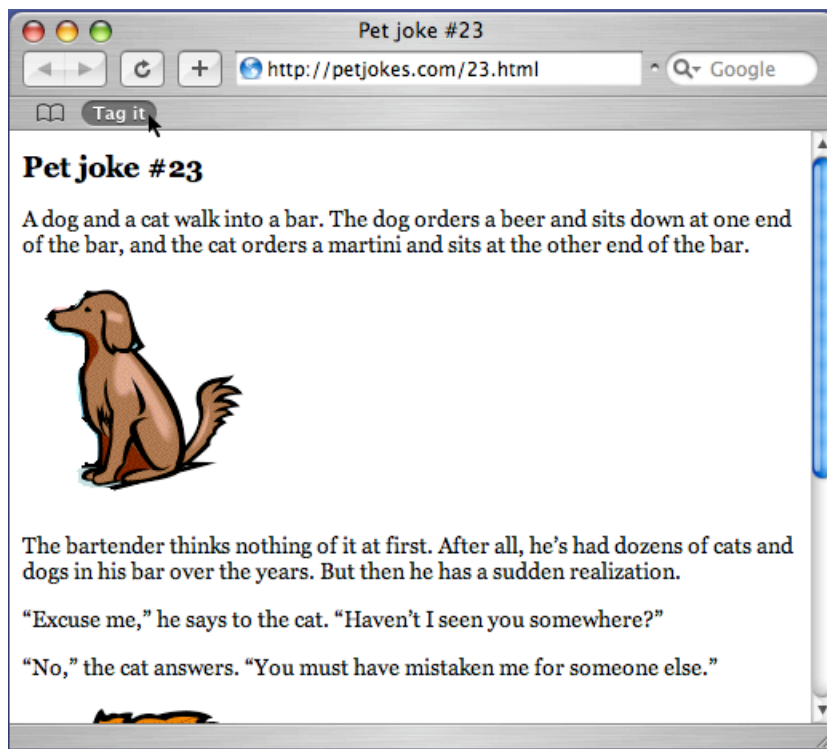


Figure 1. The resource to tag.

⁷ Most tagging services also provide alternative bookmarklets that avoid the use of pop-up windows by displaying the tag creation form within the main window.

⁸ Users sometimes 'tag' Web pages without assigning any actual tags; this practice won't be considered here.

The image shows a web form titled "tag it" in a bold, sans-serif font. In the top right corner, it says "You are logged in as Fishwick [logout]". Below this, there are four input fields: "URL:" with the value "http://petjokes.com/23.html", "Title:" with "Pet joke #23", "Tags:" with "cat, dog, funny", and "Comments:" with an empty text area.

Figure 2. The tag creation form.

Besides providing a Web interface for tagging resources and storage for users' tag records, freetagging services provide a variety of interfaces for retrieving tagged links. Typically, these include a dynamically generated HTML page for each user that shows all the tag records the user has created, another page for each tag listing all the records containing that tag, and additional pages for combinations of tags and users. Some services provide a chronological view as well, so that (for example) a user can revisit the tag records they created in a particular month or on a particular day.

Other pages typically include a front page for the freetagging site that displays the most recently created tag records, and another page with links to the tag records associated with the most commonly used tags. Most freetagging services also provide an RSS feed corresponding to each of these pages, which gives users an alternate way to discover resources associated with a tag or to follow a particular user's interests. Figures 2–7 illustrate some of these pages that users are likely to encounter in a freetagging site.⁹

⁹ These and most of the other examples given in this paper are from an imaginary site named 'Tag It' that is meant to represent a typical freetagging site.

tag it You are logged in as Fishwick [[logout](#)]

Recent tags... [[RSS](#)] Search:

July 28, 2005

Designing with Cascading Style Sheets
tagged at 13:07:44 by Eppie: [css](#), [webdesign](#) (6 other users) [[copy](#)]

Hamsters on the march!
tagged at 12:42:19 by 2e16: [hamster](#), [news](#), [toread](#) (51 other users) [[copy](#)]

Orson Welles and the Mercury Theatre
tagged at 12:40:23 by nkuitse: [orsonwelles](#), [mercurytheatre](#) (1 other user) [[copy](#)]

Painting with Pierre
tagged at 11:28:55 by pierre: [me](#) (0 other users) [[copy](#)]

Figure 3. A freetagging Web site's front page.

tag it You are logged in as Fishwick [[logout](#)]

Your links... [[RSS](#)] Search:

Dogs: inside and out
tagged May 11, 2005: [dog](#), [gross](#) (3 other users) [[edit](#)]

Petjokes #23
tagged May 8, 2005: [dog](#), [cat](#), [funny](#) (7 other users) [[edit](#)]

En recherche du chat perdu
tagged May 7, 2005: [cat](#), [sad](#) (1 other user) [[edit](#)]

Figure 4. Fichwick's saved links.

tag it You are logged in as Fishwick [logout]

Links tagged as **dog...** [RSS] Search:

Lassie: the untold story
other tags: [lassie](#), [exposé](#), [scandal](#) (2 users) [copy]

Dogs: inside and out
other tags: [anatomy](#), [showbob](#), [gross](#) (4 users) [edit]

Cave canem!
other tags: [mailmen](#), [funny](#) (8 users) [copy]

Rin Tin Tin
other tags: [rintintin](#), [story](#) (1 user) [copy]

Pet jokes #23
other tags: [cat](#), [funny](#) (8 users) [edit]

Figure 5. Tag 'dog'.

tag it You are logged in as Fishwick [logout]

Your links tagged as **cat...** [RSS] Search:

Petjokes #23
tagged May 8, 2005: [dog](#), [cat](#), [funny](#) (7 other users) [edit]

En recherche du chat perdu
tagged May 7, 2005: [cat](#), [sad](#) (1 other user) [edit]

Figure 6. Tag 'dog' as applied by user Fishwick.

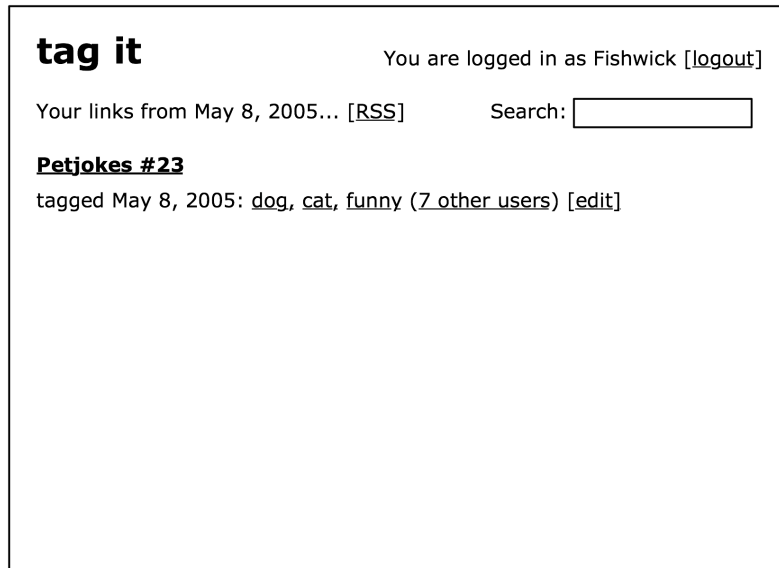


Figure 7. Tag records created by Fishwick on May 8, 2005.

Some freetagging services provide bookmarklets that take the user to a list of tag records for the Web page they're viewing.

A common visual technique used to present a set of tags is the *tag cloud*, an alphabetically ordered list of tags in which the size or color of the tags represent their frequency of use and each tag is a clickable link to a list of records associated with that tag. Tag clouds are typically used to depict the most popular tags on a site, the set of all tags employed by a particular user, or the tags associated with a particular Web page. Figure 8 shows what the tag cloud for a Web site popular with people interested in open-source programming might look like.



Figure 8. Tag cloud.

The number of freetagging sites has grown rapidly in recent months. For a long – and frequently updated – list of freetagging sites, see Free bookmark managers (2005); a chart comparing the features of nineteen freetagging services is provided in Rox (2005).

Conceptual and technical underpinnings

To more fully understand freetagging – and, especially, to see its potential uses beyond the management of personal bookmark collections – one must understand what makes it what it is. To that end, it is important first to identify the five essential properties of freetagging:

¶ It is **user-centered** – tags are devised and assigned by the users of resources rather than their creators.

- ¶ It is **uncontrolled** – there are no rules concerning which tags are allowed, what their internal structure (if any) must be, or how they may be used.
- ¶ It is **shared** – tag records are visible to everyone (except those few that are kept private).
- ¶ It is **social** – tag records are associated with the users that assigned them.
- ¶ It is **immediate** – freetagging requires minimal time, thought, and effort, and the resulting tag records are immediately available for retrieval.

Just as importantly, a successful freetagging service must be easy to use – both for users who wish to locate resources and for those who create tag records in the first place. This may seem an obvious result of the free-form nature of tagging, but it is unlikely that even the most cleanly designed freetagging service will succeed if it is slow or if it has an unwieldy interface or distracting bells and whistles.

A number of technologies come together to provide the five properties that characterize freetagging. A web server on a public web site enables sharing; a database provides storage for tag records and user accounts; a CGI or other form-handling backend program running on the server provides the interface for tag record creation and searching; client-side Javascript bookmarklets minimize the effort required to create a tag record in the first place, or to retrieve existing tag records for a resource.

Bookmarklets

The bookmarklet that the user clicks on to tag a Web page is a brief bit of JavaScript code that builds a URL to the tag creation form on the freetagging site, filling in parameters using the resource's URL and title. A typical bookmarklet might contain Javascript code such as this:

```
document.location
  = 'http://example.com/tag.cgi?url=' + escape(document.location)
  + '&title=' + escape(document.title)
```

Clicking on the bookmarklet that contains this Javascript code results in an HTTP **GET** request being sent to the Web server at `example.com`; this will trigger the execution of a CGI script on the server (`http://example.com/tag.cgi`) with two parameters set: `url` (the encoded form of the URL for the displayed resource, e.g., `http://petjokes.com/123.html`) and `title` (the encoded value of the resource's HTML `<title>` tag). The server will return an HTML page containing the tag creation form, such as the example shown in Figure 2; the URL and title fields will already have been filled in using the values that the bookmarklet provided.

Freetagging services sometimes provide bookmarklets for other purposes besides tag creation. For instance, a bookmarklet that retrieves a list of tag records for the Web page that one is currently viewing might contain Javascript code like this:

```
document.location
  = 'http://example.com/search.cgi?url=' + escape(document.location)
```

Backend code

The CGI script or other backend code that handles the HTTP request typically begins by parsing the request in order to determine what action is desired – return a tagging form, retrieve tag records, and so on – and to identify any parameters that have been supplied.¹⁰ When adding

¹⁰ Some of this work may be done by the HTTP server itself. Apache, for example, is often configured so that a user-friendly URL is automatically converted into a call to a CGI handler or

a tag record, the request will consist of the data from the completed tag creation form; when retrieving tag records, the request may consist of a user name, one or more tags, or some other combination of parameters. The backend code passes this data to a separate storage mechanism, which – depending on what action has been requested – creates a record or fetches the desired records from disk (or from an in-memory cache). (If the user is simply requesting a tag creation form, this step is skipped.) In any case, the final step will be to construct an appropriate HTML page and return it to the user’s browser.

Storage

A relational database is the most common type of storage used in a freetagging service; such a database allows for a clean, formal design and will scale much better than simpler forms of storage such as tab-delimited or other flat-file text databases. (The ability to handle large numbers of users and tags is important because the usefulness of a freetagging service is generally limited by the number of users and the number and variety of tag records created.)

Data models

Although freetagging implementations vary in the types of data that may be included in a tag record, three data elements are always present: *user* (a user-selected ‘handle’ or nickname), *resource* (a URL or, in the case of Flickr and similar sites, the contents of the resource itself), and *tags* (zero or more text strings). These common data elements, and the ways in which they are related, form the center of freetagging data models.

Other data elements in tag records provided by some freetagging services include free-text comments, resource ratings (e.g., from 1 to 5), user-supplied excerpts from the tagged item, and the URL of the referrer (the Web page – if any – from which the user arrived at the tagged resource).

For details on common freetagging data models, see Keller (2005).

Alternate interfaces

Keeping in line with the simplicity of tagging, most freetagging Web sites follow simple conventions for constructing the URLs that lead to the various dynamically generated pages described above. For example, in unalog¹¹ the list of all records containing the tag ‘cat’ may be retrieved using the URL `http://unalog.com/tag/cat` and the list of records tagged by the user ‘moe’ using `http://unalog.com/person/moe`. These user and tag elements may be combined in a single URL – `http://unalog.com/person/moe/tag/cat` – to retrieve only those records containing the tag ‘cat’ that were created by the user ‘moe’. Tags may be combined using a plus sign (+), which functions as a Boolean AND operator, and adding /`rss` results in a link to an RSS feed of tag records instead of an HTML page; hence the URL `http://unalog.com/tag/cat+dog/rss` will retrieve an RSS feed of all records containing the tags ‘cat’ and ‘dog’.

Conventions such as these are useful for Web-savvy users, making it possible, in essence, to construct certain Boolean tag searches using the browser’s location field as a sort of search command line. Web authors benefit, too, by gaining the ability to link to a dynamic (or fixed) set

other backend code – for instance, the URL `http://example.com/tag/cat+dog` might be translated to `http://example.com/retrieve.cgi?tags=cat+dog`.

¹¹ `http://unalog.com/`

of resources without having to create a separate page for them. (It remains to be seen whether ordinary users will make use of such URL patterns.)

Many freetagging services go a step further and provide a Web application programming interface (API) that others can use to create their own tag-aware tools. These APIs go beyond the simple conventions of tag retrieval URLs, providing freetagging functions such as tag record creation that would otherwise require human intervention. For example, Biddulph (2004) presents Python code that uses the del.icio.us API to support the following retrieval tasks: (1) identify the tags with which a resource has been tagged, and by whom; (2) find resources that have been tagged using the same tags as a given resource; (3) given a set of tags, determine the resources to which they have been applied and find all other tags that have been applied to the same resources.

Together with the simple tag building conventions, these Web APIs make it possible to automate the creation, modification, and retrieval of tag records and user information, providing a means for Web programmers to build their own value-added services on top of the capabilities provided directly to users by freetagging sites. A number of such services have already come into being: for example, Doctorow (2005) describes the tag capabilities of Technorati,¹² a weblog aggregation site that uses the APIs and regular URLs of other freetagging services to display tag records from those sites as well as its own.

Freetagging origins

When considering the origins of freetagging and of today's freetagging services, it is important to examine not only their ancestry but also the conditions that led to their creation and under which they now thrive. These conditions include:

¶ *The proliferation of web applications.* Exposure to increasing numbers of helpful, everyday services such as Mapquest and eBay has created an expectation among users that that's the way the Web should be.

¶ *The rise of the user as author.* As Web users have grown accustomed to discussion lists, Amazon reviews, wikis, easily maintained blogs, and so on, they have come to expect to be able to contribute to the Web, not just to consume it.

¶ *The rise of full-text search engines.* As the Web has grown beyond anyone's ability to organize more than a tiny fraction of it, hierarchically organized Web directories such as Yahoo! have largely given way in users' consciousness to the ever-increasing need – and ability – to find (enough) things using unformalized, unmediated¹³ keyword searches. (For a provocative view of the decline of hierarchical organization on the Web, see Hiler, 2005.)

¶ *The proliferation of heterogeneous systems on the Web.* When the Web was read-only (at least for ordinary users), it was more homogeneous; the explosion of disparate writeable systems in recent years – for example, blogs, web applications, and wikis – created more and more barriers to interoperability.

Freetagging *per se* appears to have begun in late 2003 with the birth of the del.icio.us Web site, but there were a number of precursors. The most immediate ancestors of freetagging are weblogs and bookmarking tools; other software with some claim to ancestry includes wikis and web-based discussion boards.

¹² <http://www.technorati.com/>

¹³ In the sense of 'immediate, unfiltered, not the product of artifice.'

An early glimpse of some of the thoughts behind freetagging can be found in Merholz (2000), which argues against “the tyranny of hierarchy” and proposes an alternative with three essential properties:

Is there a bottom-up solution?

- Basic-level categories – most immediately understandable to human perception
- Assign attributes (metadata) and use that to create relationships
- Allow people to personalize the information so it’s relevant to them

The remainder of this section presents a variety of technologies that contributed to the development of freetagging, comparing some of their uses with those of the latter.

HTML-based metadata

The `<meta>` element provided in the HTML standard makes it possible for Web authors to include simple metadata in their documents; a `<meta>` element whose ‘name’ attribute is set to ‘keywords’ is commonly used to hold author-assigned descriptors – for example, ‘pets’ or ‘nuclear energy’ — that indicate the page’s subject matter. Like the tags assigned by users of del.icio.us and other freetagging services, these keywords may take any form (subject to the usual restrictions placed by the HTML standard on the values of attributes). The use of author-assigned free-text keywords goes back much than this, of course, but the keywords commonly placed in HTML metadata elements can be seen as an important ancestor to the free-text tags used in freetagging.

Bookmarking tools

The first graphical bookmarking tool was the ‘hotlist’ function of NCSA Mosaic, first released in 1993 (National Center for Supercomputing Applications, 2002). Later browsers added improved bookmarking capabilities, including the ability to organize bookmarks hierarchically in folders and subfolders. As users began to use more than one browser on more than one computer, the need for centralized bookmarking arose; the personal bookmarking tools that resulted were only a step away from *social* bookmarking tools.

Weblogs

To understand the role of weblogs in the birth of freetagging, it is helpful to first divide them into two major types: *journals* and *linkblogs*. (The latter are sometimes called *filters*; see Caslon Analytics (2005) for a similar analysis.) Entries in a linkblog consist of a link to a web resource and some commentary, generally quite brief; journal blogs inform the reader of the blogger’s everyday experiences and tend to fall more along the lines of ‘Today I had a cheeseburger for breakfast and then met up with my friends.’ Though the latter would seem to dominate blogging today, they don’t have any real relevance to freetagging.

Over time, blogging software came to allow the blog author to specify the category of each posting; today this is true of most blogging software.¹⁴ Category names are generally limited only by the blogger’s imagination (though special characters and spaces are often not permitted); the ease with which they can be created varies, but typically only a minor effort is required. (In *blosxom*, for instance, creating a category entails creating a directory on the server.)

In some respects, the blogger’s use of categories parallels the tagger’s use of tags. Although freetagging requires no effort at all to create a tag – a fact with obvious positive and negative

¹⁴ For example, *blosxom* – <http://www.blosxom.com/> [check URL]

consequences – both blog entries and tag records are centered around a link to a Web resource and have a social aspect: not only does a blog entry indicate the person who wrote it, but the practices of allowing signed comments and ‘trackback’ links draw relationships among bloggers. As was true of early bookmarking Web sites, however, there is a more significant difference: weblog entries, unlike tag records, may generally be labeled using only a single category.

Wikis

Leuf and Cunningham (2001, p. 16) describe the essential properties of the Wiki concept:

At the functional level, which is what the user sees, the essence of Wiki can be summarized by these statements.

- A wiki invites all users to edit any page or to create new pages within the wiki Web site, using only a plain-vanilla Web browser without any extra add-ons.
- Wiki promotes meaningful topic associations between different pages by making page link creation almost intuitively easy and by showing whether an intended target page exists or not.
- A wiki is not a carefully crafted site for casual visitors. Instead, it seeks to involve the user in an ongoing process of creation and collaboration that constantly changes the Web site landscape.

Wiki is a lot about a collaboration space, albeit an unusual one because of its total freedom, ease of access and use, simple and uniform navigational conventions, and apparent lack of formal structure. Wiki is also a way to organize and cross-link *knowledge*...

The same statements can be applied, with only minor qualification, to freetagging sites: both provide users with a deceptively simple foundation on which to build information environments incrementally using iterative, shared changes. Neither is intrinsically collaborative – users can choose to work together with common goals, or they can choose not to – but both make a sharing of effort not only possible but easy while not sacrificing individual users’ ability to use these tools for their own purposes and in a way that makes sense to them personally.

One aspect of wikis is particularly relevant to a discussion of freetagging. As is true of the Web and of information systems generally, a growing wiki needs pruning: the more content a wiki contains, the more important it becomes to organize that content. Contributors to early wiki efforts tried a variety of different approaches to filling this need, including the ‘guide’ page – a single page consisting of links to a set of related pages within the wiki that might usefully be considered parts of a whole, with some expository text to help the user find their way among those pages. Creating a guide page required a good deal of thought, but could be built incrementally just like any other wiki page, and the ease with which links are created in a wiki made the job at least technically simple; guide pages continue to play an important role in wikis today.

Another early innovation was to automatically include a *backlink* in each wiki page – a button or hypertext link leading to a dynamically generated list of all the other pages in the wiki that link to the original page. (Some wiki implementations automatically turn the page’s title into a backlink.)

Some users, looking for a way to erect signposts in the jumbled networks of wiki pages, hit upon the idea of adding category labels at the ends of pages; because of the primacy of links in a wiki, it was only natural to create these category labels in the form of links, and these special links soon came to be known as *badges*. Figure 9 shows a wiki page with two badges – ‘CategoryOpenSource’ and ‘CategorySocialSoftware’.

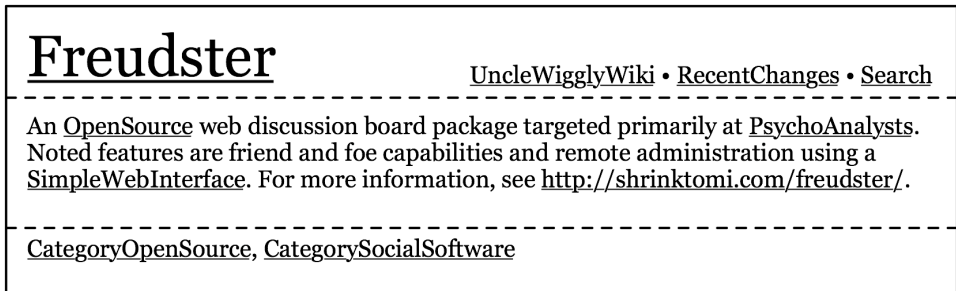


Figure 9. A wiki page with badges.

When viewing a page such as this, a user can retrieve a list of all pages bearing the same badge by first clicking on the badge (for example, ‘CategoryOpenSource’), which takes the user to the page of that name. As Figure 10 shows, this page may contain a description of the badge’s intended use.

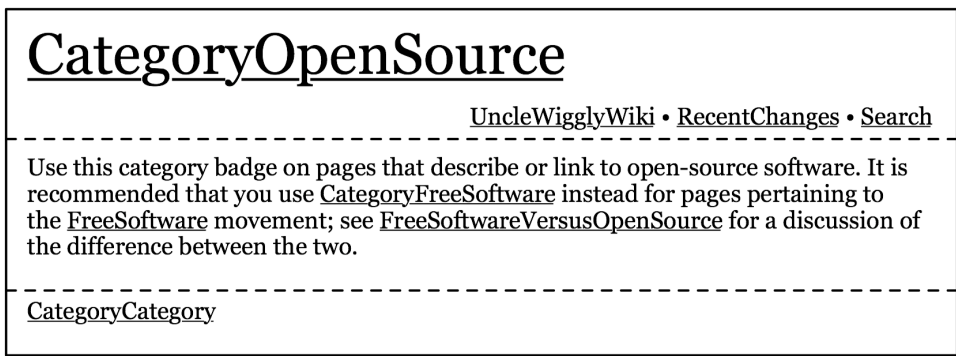


Figure 10. The page describing a wiki category badge.

Once there, the user clicks on the badge page’s backlink – in this case, the page’s title serves that purpose – to recall the dynamically generated list of pages containing that badge.

Wiki badges have a wide variety of uses, indicating not only a page’s topic (‘CategoryWebServices’, ‘CategorySocialSoftware’) but its status (‘CategoryDeleteThisPage’, ‘CategoryRefactor’) or type (‘CategoryDiscussion’, ‘CategoryCategory’). They are strikingly like the uncontrolled, undifferentiated, multipurpose tags in freetagging systems, but they are at least potentially controllable since any user can edit a page’s badges. At their core, both practices – tagging resources and applying wiki badges – provide user-created labels that can be employed for any purpose. The essential differences are that freetagging systems provide the ability to label *any* resource – not just the pages of a single Web site (the wiki) – and wiki badges are represented by pages that are user-editable and so may be used to describe the badge’s function or relate it to other badges – unlike the pages in a freetagging site, such as are depicted in Figure 5, which users may not edit.

Social software

Examples of software with an essential social component are easy to find: they range from e-mail, USENET, chat and instant messaging to videoconferencing, Web discussion groups, and relatively recent developments such as Friendster and Orkut. Although these applications represent a wide variety of purposes, they all provide a means by which users ‘speak’ and ‘listen’ to each other and a forum in which to do so. Furthermore, they share one absolutely essential property: they provide a mechanism – in the form, say, of an e-mail address or user-selected

'handle' – by which every user is given a *voice* of her own. User voices are integral to freetagging because every tag record imparts the view *of a particular user*, and a tag that means one thing to one user may mean something quite different to another.

Applications of freetagging

Because users are given free rein to tag Web resources in whatever ways further their own particular needs, the uses of freetagging services are potentially as various as the people that use them. However, at this early stage a number of overall usage patterns are clear. (The next section will address the specific applications of tags.)

Bookmarking

Freetagging is commonly used as a bookmarking tool – or, as Fichter (2004) describes succinctly in the title of her article, “A tool for finding things again.” Users of freetagging services save URLs along with their chosen tags precisely so that they – and perhaps others – can more easily retrieve the resources later. Roush (2005) sees this as the main purpose of freetagging:

The main point of [freetagging] sites is to allow members to store the addresses of Web pages that interest them, which they can later revisit from any computer or browser. Along the way, they are invited to add tags to the pages they bookmark – descriptive terms of their own choosing that make it easier to find favorite sites in a big stack of bookmarks.

Organizing

Like the bookmarking interfaces of most Web browsers, these services may be used not only to store a user's links, but also to organize them in whatever way makes their subsequent retrieval easiest for the user. In a Web browser, this might involve the use of category-based folders and subfolders into which links are placed; in freetagging, the organization is accomplished by means of tags. Of course, tags aren't hierarchical – indeed, as Mathes (2004) and others point out, the lack of hierarchical relations is an essential feature of tagging systems – but this limitation is mitigated by the ease in which a single bookmark may be placed in any number of categories. Lund (2005) argues that the unstructured 'clutter' of a freetagging system is well suited to user needs:

The 'tags' of Connotea and other social bookmarking tools often lead commentators to decry the anarchy of unconstrained keywords, but this overlooks the fact that tags are intended first and foremost as a way for individual users to manage their own collections. In this way, they are similar in purpose to folder names in computer file systems. However, instead of creating sub-levels of organisation by nesting folders hierarchically, flat tagging achieves this by assigning multiple tags to each item, each tag being treated equally. This releases the user from some of the constraints of a traditional file system...

Sharing

The bookmarking aspect of freetagging services is useful enough that some users will never go beyond this, but the added social dimension of freetagging promises much more. Gottlieb and Dilevko (2001) describe the benefits that a shared bookmarking system brings to an organization, as well as the challenges that such a system entails:

As a pooled resource, a form of bookmark database could provide collective access to information of use to the members of an organization, thus benefitting the organization as a whole. ... The difficulty in organizing bookmarks for multiple users is that members of the intended audience would have different conceptions of or ways of defining the web sites...

Fichter (2004) likens freetaggers to foraging ants, leaving pheromone-like tags to mark the objects of interest that they encounter:

[Users] act independently, gathering what interests them. Collectively, their actions create a Web of networked sites, information, and people. The social bookmarking applications analyze this network of information, links, and people to recommend new sites and display linking behavior patterns that help [users] track trends and discover new information.

Continuing the ant metaphor for a moment, from the succession of individual acts of tagging – imprecise though they may be – collective ‘paths’ of sorts emerge; for example, each time a particular Web resource is tagged ‘cats’ in a freetagging service, the likelihood increases that users of the service who are interested in cats will find it. What’s more, by increasing the tag’s frequency, this tends to reinforce the use of the tag ‘cats’ over that of other tags – for example, ‘felines’ – and the repeated application of the ‘cats’ tag to the Web resource helps to shape the collective view (‘this is about cats’) of that resource.

Freetagging in context

Armed with an understanding of freetagging, its conceptual and technological underpinnings, its origins, and its uses, it may now be possible to place it in a larger context. With that goal in mind, we will now compare freetagging with conventional metadata – in particular, subject-oriented metadata – and explore two parallels with the vision of the Semantic Web.

Metadata

Metadata has been widely – and rather loosely – defined as ‘data about data’ (Schwartz, 2001, p. 8). The kinds of metadata that have been the subject of most efforts to date are primarily those provided by the authors of works or by trained professionals such as catalogers and indexers. The creation of metadata by software using latent semantic indexing or other automated means holds promise but is beyond the scope of this paper.

Metadata practices today range from the relative anarchy of HTML <meta> tags to the bare-bones, easy-to-use metadata elements of the Dublin Core (Weibel, Kunze, Lagoze, and Wolf, 1998) to the highly structured Library of Congress subject headings (Library of Congress Cataloging Policy and Support Office, 2005).

The tags seen in freetagging sites are also a form of metadata; unlike conventional metadata, however, tags are chosen and applied in a free-form manner by *users* without any rules beyond such minimal restrictions as the forbidding of commas in a tag. A comparison of the two approaches must address three concerns: the kinds of data elements that each provides, the ways in which the data elements are applied, and the data structures that result.

Data elements

A conventional metadata system – even a minimalist system such as Dublin Core – provides a number of data elements that may be used in a metadata record, each of which denotes a specific aspect of the item that the record describes – its *author*, the *title* by which it is known, its *subject matter* (say, Web sites or Teddy bears), what it *is* (for example, a Web site or a Teddy bear), and so on. In many metadata environments, data elements that denote an item’s subject matter may be further divided, for example, into personal, geographical, or topical subject headings.

A typical freetagging service, on the other hand, provides only a few fields for users to fill out when creating a tag record – URL, title, tags, and comments – and only the first two have a specific intended meaning. There is no intrinsic reason not to provide additional fields for tag

records; indeed, some sites – most notably, CiteULike¹⁵ – offer many more. But the focus of freetagging, of course, is on the tags themselves, and so this must be the focus of our comparison of conventional metadata practices with freetagging.

Tag uses

In a study of tags found on two freetagging services' web sites, Mathes (2004) found that users most commonly used tags to indicate subject matter, but that other applications were also common. A variety of tags found at del.icio.us indicated genre (for example, 'comics'), medium ('photography'), or the user's feelings toward the tagged resource (for example, 'wishlist' was apparently used by many users to tag desired consumer products); at the photo-sharing site Flickr, colors were commonly used as tags, as were years. The types of associations that users of freetagging services most commonly make when assigning tags to Web resources can be seen to include the following (tags set in italics):

- ¶ Topic – 'It's about a *hamster*.'
- ¶ Form – 'It's a *comic strip*.'
- ¶ Location – 'It's set in a *sauna*.'
- ¶ Descriptive characteristic – 'It's *funny*.'
- ¶ Relevance – 'It's *important*.'
- ¶ Application – '*Show Janet*.'

While the first three of these are commonly represented in conventional metadata records, the last three are decidedly not. As this shows, what distinguishes freetagging from other metadata systems is the flexibility not only in the *forms* that tags may take, but the flexibility in the *purposes* to which tags may be put. The lack of rules or guidelines frees the person *tagging* a resource from the chore of identifying the aspect denoted by a tag, leaving it to the person *retrieving* a tag record to do so. A similar compromise is made in the post-coordinate subject vocabularies of periodicals indexes, in which any given subject descriptor might be a personal name, a topic, an era, and so on.

Naturally, the result may well be ambiguity, due to the lack of a controlled vocabulary – or, indeed, of any control at all. Examples of misspellings and misapplication of tags abound in freetagging sites, but a more interesting – and more subtle – ambiguity is illustrated by Hicks (2005) using an example from the photo-sharing site Flickr:

A user who travels to Tokyo might tag all photos from the trip as "Tokyo," including those taken inside a hotel room... But other users might expect to see only photos of the Tokyo cityscape, and not a hotel's interior, when viewing photos tagged as Tokyo.

In some cases, it may be possible to make probabilistic guesses as to which aspect of a resource is being denoted. For example, as Coates (2005) points out, the tag *blogs* – plural – is likely to indicate that a Web page is *about* blogs, whereas the singular form *blog* is more likely to signify that it *is* a blog. In general, however, this kind of determination cannot be made with even a minimum of certainty.

It is possible – at least in principle – to use tags to specify precisely the same data elements as are provided conventional metadata data elements – author, title, and so on. Indeed, some users have attempted to do just this, assigning tags such as 'author:shirky' or 'location:tokyo'. This type of tagging is very uncommon, however, and, in any event, hardly fits in with the simple, immediate nature of freetagging generally.

¹⁵ <http://www.citeulike.org/>

Capturing the thought bubble

As Quintarelli (2005) points out, formal systems of categorization and classification are best suited to relatively stable, static environments – conditions that do not readily apply to the Web:

To classify cleanly, formal categories should be identified that do not change over time, contain homogeneous entities and are capable of describing all the items in a corpus. ...

Using a sound and complete classification scheme requires professionals to do the job, a common clear view of the domain and skilled users that understand the categories and the structure of the classification to use it without problems.

On the other hand, sprawling, heterogeneous information sources make up an enormous, ever-changing, time-sensitive, not-clearly defined corpus of items to classify without a central authority, targeted at a heterogeneous and increasing group of users. This situation requires **new and different classification strategies**.

The Web today fits neatly in this description.

Perhaps the most commonly cited advantage of freetagging over the metadata practices commonly undertaken by indexers and catalogers is that the use of free-form tags to categorize or otherwise label resources requires no formal training, because no controlled vocabulary or syntactical rules are used – tags may take any form and may relate to a resource in any way. On the contrary, the cognitive function of ‘tagging’ – assigning labels and categories – is an innate human trait.

The creation of conventional subject metadata involves a relatively great deal of mediation: see a resource; figure out what it’s about; choose the appropriate form for the subject metadata; enter it in the appropriate encoding; repeat. Freetagging, on the other hand, comes much closer to capturing the user’s immediate thoughts: see a resource; enter words or phrases that it makes you think of; repeat.

In this way, freetagging imposes only a minimal tariff on recording associations with a resource (what it’s about, what it is, who wrote it, what to do with it). This relative lack of mediation in freetagging makes it possible for any Internet user to tag resources, and users are able to characterize many times the number of resources that conventional metadata creation allows.

Figure 11 illustrates – in a crude way – the thought processes of three users looking at the same Web page (the page shown in Figure 1), which contains a joke (“A dog and a cat walk into a bar...”) and images of a dog and a cat. The thought bubbles depict labels that the users might assign to the page, and the users are arranged roughly along an axis of immediacy – in other words, how near the labels are to the thoughts that pop into their heads when looking at the page. At the high end of the mediation spectrum, Yolanda appears to have thought of two Library of Congress subject headings and Otho seems to be thinking in terms of facets; closer to the other end, Eppie and Fishwick have less formalized, more immediate labels in mind.

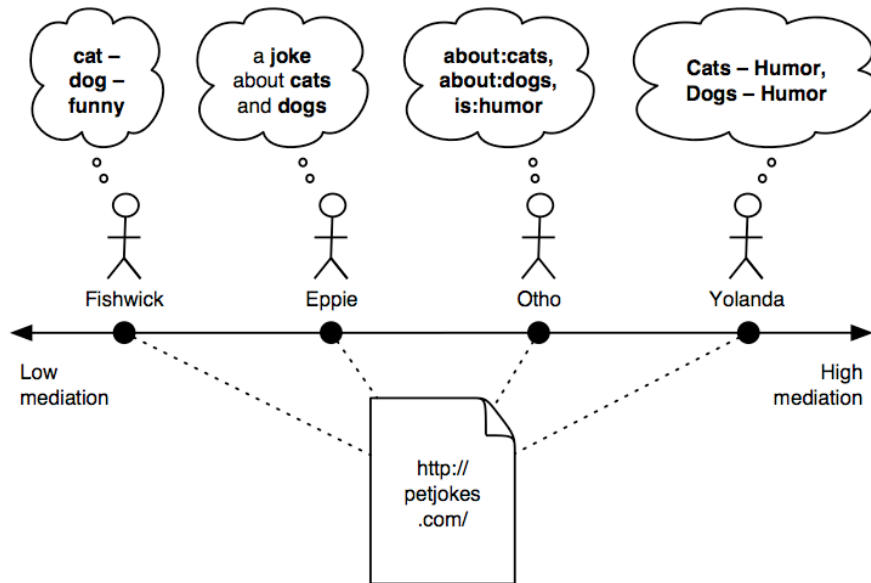


Figure 11. Capturing the thought-bubble.

While most efforts to categorize the Web have focused on users like Otho and Yolanda who have the special skills – and time – required to take advantage of formalisms such as LC subject headings or ERIC descriptors, freetagging exploits the immediate, unformulated (or ‘unsophisticated’) thoughts of users like Fishwick and Eppie.

Freetagging and the Semantic Web

The Semantic Web is a vision, first articulated by the inventor of the Web, Tim Berners-Lee, of a future Web in which widespread metadata conveys the meaning of Web resources and their component parts in a way that is comprehensible to software tools, not just to people:

The Semantic Web will bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users. (Berners-Lee, Hendler, and Lassila, 2001).

For some Web authors, freetagging is the antithesis of the Semantic Web, which they see as a ‘pie-in-the-sky’ dream that ignores the immediate needs of Web users – see, for example, Shirky (2003). But the two might not be so far apart, despite their differences – most notably, the stark differences between the heavy formalisms of the Semantic Web and the uncontrolled free-for-all of tagging.

Ontologies and folksonomies

Parsia (2003) emphasizes the potential of the Semantic Web to integrate disparate resources and services on the Web, contrasting this vision with the current state of the Web:

If the original Web is a hypermedia system with aspirations toward KR [*i.e.*, *knowledge representation*], the Semantic Web seeks to be a KR system deeply integrated with global, distributed multimedia. ... [In today’s Web] there are no content-directed constraints on putting up a Web page or with making a Web link to anything else on the Web. These design constraints are commonly expressed with the slogan, “Anyone can say anything about anything.” ... The difference between the current Web and the Semantic Web is that the

Semantic Web asks people to say their anythings about anything in a way that's more amenable to significant, content-sensitive machine processing.

The power of the Semantic Web lies in the inferences that can be drawn from these “anythings about anything” – statements about the resources of the Web and the relationships among them. To make these inferences, such statements must be embodied in a form that is readily ‘understood’ by software – in other words, as metadata. This metadata takes the form of *triples*; a triple is a statement about a resource (1, the *subject*) that denotes its relationship (2, a *verb* or *attribute*) with another resource (3, the verb’s *object* or the *value* of the attribute). The resources and relationships are all specified using URIs; for example, a statement that “Ulysses Fishwick is the author of *One hundred poems about rutabagas*” might use the URI <http://www.fishwick.com/~ulysses/> to represent Ulysses Fishwick, <http://relations.r.us/writtenby/> to denote the author relationship, and <http://poetaster.org/fishwick/rutabagas.html> to represent the written work. Figure 12 illustrates this statement.

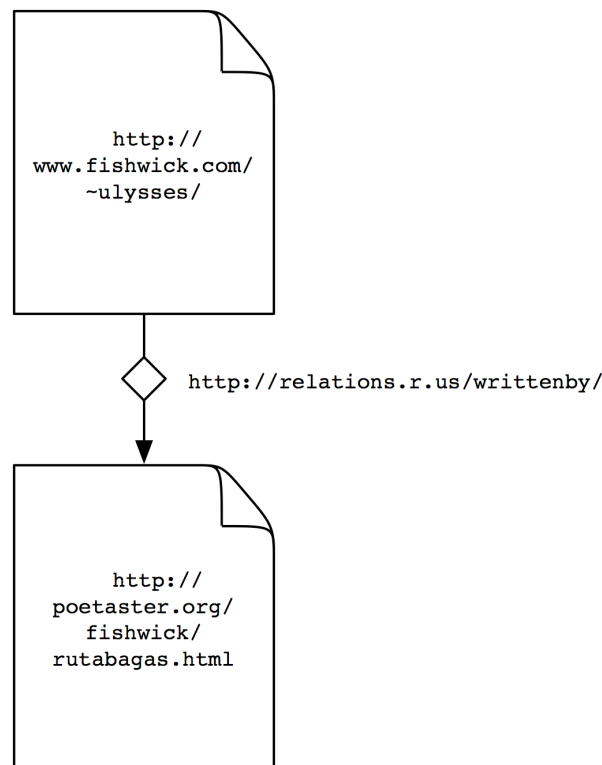


Figure 12. A Semantic Web triple.

In the Semantic Web, such statements are encoded in a formal language and collected in an ontology, “a partial conceptualization of a given knowledge domain, shared by a community of users, that has been defined in a machine-processable language for the explicit purpose of sharing semantic information across automated systems” (Jacob, 2003). The Semantic Web allows for the existence of any number of ontologies; indeed, ontologies may be explicitly related to each other.

In late 2004, Thomas Vander Wal coined the term *folksonomy* to denote the set of tags used in a freetagging site, seen as a whole (Smith, 2004), drawing an ironic parallel with *taxonomy*. This

neologism has caught on and is now in widespread use. (It is also used to denote the phenomenon of freetagging more generally.)

Folksonomies have none of the structure of such formalisms as ontologies, taxonomies, and thesauri, which consist not only of concepts or classes – for instance, MAMMAL and DOG – but also of statements of the relations between them – for example, “DOG **is a** MAMMAL” or “**mammal:** NT dog.” Many freetagging services, however, attempt to draw connections between tags by looking for different tags that have all been associated to the same resource; the more often two tags are associated with the same resource, the more they are said to be related. These connections between tags are then used to enhance the display of tag records, giving users further avenues for discovering resources and suggesting tags that users might wish to use instead of – or in addition to – the tags they had in mind when creating a tag record. Similar relations can be drawn between users, for example to refer one to those other users whose tag records might be of interest to them. Implementers or freetagging sites are only just beginning to scratch the surface of these possibilities.

A clear parallel can be drawn between the statements in ontologies that form the heart of the Semantic Web and the tag records found in the folksonomies produced using freetagging services. Freetagging produces triples of its own: a tag record is a statement by a user (1) that associates a tag (2) with a Web resource (3). What’s more, despite the informality of freetagging, the regular URL ‘syntaxes’ – described above – that many freetagging sites provide make it possible to capture the ‘statements’ in a folksonomy using a formal language, much as occurs in the Semantic Web. For example, if the Web site for Apple Computer’s iPod is tagged ‘ipod’ by the user ‘moe’ within the Unalog freetagging service, the resulting tag record can be encoded using three URIs: <http://unalog.com/person/moe> represents the user, <http://unalog.com/tag/ipod> denotes the tag, and <http://www.apple.com/ipod/> represents the iPod. Figure 13 provides an illustration of this.

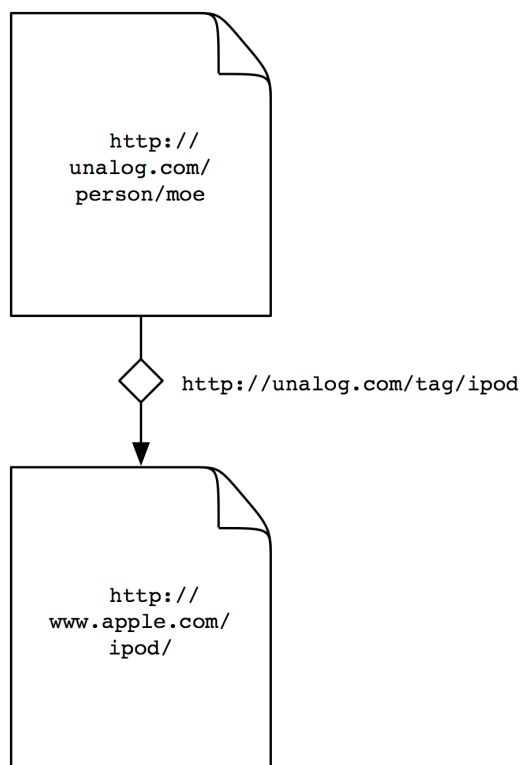


Figure 13. A freetagging ‘triple’.

There are important differences to note, however. First, such statements in the Semantic Web specify the *nature* of the relation between the subject and its object: whether the subject is the author of the object, whether the subject is a subtype of the object, and so on (there is no limit to the relationships that may be drawn). As we’ve seen, this specificity is lacking in freetagging. Second, freetagging brings *users* into the picture: a folksonomy’s ‘statements’ are always subject to the qualification *according to user x*, and to all the good and bad that that implies.

The point of this comparison is not that freetagging is a suitable substitute for the Semantic Web, but rather that the folksonomies produced by freetagging – even today’s early incarnations – can be represented formally, and that this ability to make machine-processable statements about tag records raises a number of interesting questions. For example, how might formalized folksonomies be linked to the ontologies of the Semantic Web? Is it possible to relate the tags used in one freetagging site with the tags used in another? And how might the centrality of users, which is such a key precept of freetagging, be brought to the Semantic Web?

Annotea

Another interesting parallel with the Semantic Web can be found in Annotea, an experimental project of the World Wide Consortium (W3C)’s Semantic Web initiative that anticipated many of the features of today’s freetagging services. In an introduction to Annotea, Koivunen (2001) describes it as a client-server combination that promotes collaboration using shared bookmarks with user-supplied annotations; these annotations may be “comments, notes, explanations, or other types of external remarks” associated with a Web resource:

When the user gets the document he or she can also load the annotations attached to it from a selected annotation server or several servers and see what his peer group thinks. Similarly

shared bookmarks can be attached to Web documents to help organize them under different topics, to easily find them later, to help find related material and to collaboratively filter bookmarked material.

As part of this experiment in Web annotation, the W3C implemented an Annotea server and added the ability to create and view annotations to its Java-based browser, Amaya. Later, other developers created Annozilla¹⁶, a plug-in for Mozilla that provides annotation display – but not creation – capabilities. Annotea was an experimental project and as such is no longer active, but Annozilla is still actively maintained and the W3C continues to provide a working Annotea server at <http://annotest.w3.org/annotations>. For a more detailed description of Annotea, see Koivunen, Swick, and Prud'hommeaux (2003).

The future of freetagging

Some freetagging systems have branched out into other functions besides the simple tagging model. For example, CiteULike and Connotea¹⁷ are online reference management services with tagging capabilities that enable users to build and manage collections of links to e-journal articles and other online resources.

Freetagging (of a sort) is also being used by authors (typically of blogs) to present topical metadata; this use goes back to the keywords used in blog postings before freetagging existed. Gibson's (2005) xFolk proposal is a good example of the work being done in this area, as is the unfortunately named but fully functional tink¹⁸.

Increasingly, tagging capabilities are being added to Web services and other applications that don't strictly fit the definition of a freetagging service. For instance, Yahoo! recently announced their MyWeb 2.0 effort, which provides tagging services – but, strangely, not the ability search using a combination of free text and tags – to its users.

Some authors have envisioned ways of bringing freetagging closer to controlled vocabularies, either by changing the nature of freetagging or by adding an 'expert layer' on top of it. An example of the former is given by Wright (2004), who imagines:

... a scenario where end-users could freely create, adopt or reject terms stored in a distributed repository that gets administered by a representative authority that "owns" the vocabulary: normalizing terms, identifying and mapping semantic relationships (determining synonyms, related terms, parent-child relationships and such).

On the other hand, Wright's 'end-users' might in fact turn out to be the traditional purveyors of metadata, precisely those trained professionals that are currently struggling to keep up with the skyrocketing amount of organizing to be done: catalogers, indexers, and Internet architects. These experts might be able to use tags as a sort of raw natural resource, drawing connections – if, perhaps, only probabilistically – between tags and formal systems. For example, a subject gateway might link subject descriptors and tags in such a way that a metadata record containing the descriptor 'Motion picture' could be displayed with links to the records associated with the tags 'movie', 'film', and 'cinema' in a particular freetagging service (or on several of them).

The future of freetagging may be less bright than some suppose. One concern includes the prospect of frequently-used tags becoming meaningless for retrieval purposes, which Rosenfeld (2005) raises:

¹⁶ [Need URL]

¹⁷ <http://www.connotea.org/>

¹⁸ <http://tinky.org/>

I also wonder how well Flickr, del.icio.us, and other folksonomy dependent sites will scale as content volume gets out of hand. Even now, for example, uploading your summer vacation photos to Flickr and tagging them “summer” will group them with over 6,000 other photos. Hard to browse now, harder when there are 60,000 photos a year from now. And it’s a safe bet that no one will bother to go back and re-tag their photos with more precise terms.

Some subject heading systems, such as the LC subject headings, avoid this hazard by mandating the use of the most specific subject headings applicable to a resource. On the other hand, in the shifting tides of the Web – and of human endeavor generally – new specificities may arise just as new concepts do (though, one hopes, not as often).

Perhaps the greatest danger to the future of freetagging, however, is a familiar one: spam. Intrusive, off-topic, repetitive, and offensive postings are often said to have ruined USENET; some would say that at this point spam has ruined e-mail too. If freetagging continues to rise in visibility and importance, its appeal to spammers will likely rise apace. This is an important area that merits further investigation.

Conclusion

Freetagging services have proliferated rapidly in the months since their inception, and there’s little doubt that more will spring up in the near future. Some of these will come and go, of course, but the deep roots of freetagging – and the rapidly increasing use of user-selected tags in an increasingly user-centered Web – indicate that tags are here to stay.

To ignore either the promise or the perils of freetagging and the folksonomies that result would be a lost opportunity to shape this new technology. While the appeal of freetagging is clear, the full measure of its usefulness – and of its limitations – has not been taken; scholarly studies of tagging practices and a fuller analysis of the folksonomies that result would constitute an important first step in the right direction.

Of all the things that have been said and written about the subject, one sentiment in particular seems to best capture both the promise of freetagging and the horrors that many associate with it: that freetagging exploits the raw power of the Internet masses to find a brute force solution – imperfect and uncontrolled – to the problem of organizing information. As library professionals, we should pay particularly close attention to the possibility of harnessing this raw energy for our own users, and of exploring ways to apply our hard-earned understanding of categorization, classification, and metadata to this exciting – and daunting – new phenomenon; at the very least, we may be able to learn some important lessons from the experience.

References

- Bates, Mary Ellen. (2005). Furl it! *Online*, May/June 2005, 64.
- Berners-Lee, Tim; Hender, James; and Lassila, Ora. The Semantic Web. *Scientific American*, 284(5), 32–43.
- Biddulph, Matt. (2004). *Introducing del.icio.us*. Retrieved July 17, 2005, from <http://www.xml.com/pub/a/2004/11/10/delicious.html>
- Caslon Analytics. (2005). *Blog profiles: blogging*. Retrieved July 17, 2005, from <http://www.caslon.com.au/weblogprofile1.htm>
- Coates, Tom. (2005). *Two cultures of fauxnomies collide*. Retrieved July 15, 2005, from http://www.plasticbag.org/archives/2005/06/two_cultures_of_fauxnomies_collide.shtml

- Doctorow, Cory. (2005). *Technorati tags: three great services on one page*. Retrieved July 28, 2005, from http://www.boingboing.net/2005/01/13/technorati_tags_thre.html
- Fichter, Darlene. (2004). Tools for finding things again. *Online*, 28(5), 52–56.
- Free bookmark managers. Retrieved July 15, 2005, from <http://www.lights.com/pickalink/bookmarks/>
- Gibson, Bud. (2005). *xFolk Entry 0.4 – microformat for decentralized tagging*. Retrieved June 4, 2005 from http://thecommunityengine.com/home/archives/2005/05/xfolk_entry_04.html
- Gottlieb, Lisa and Dilevko, Juris. (2001). User preferences in the classification of electronic bookmarks: implications for a shared system. *Journal of the American Society for Information Science and Technology*, 52(7), 517–535.
- Green, Heather and Hof, Robert D. (2005). Picking up where search leaves off. *Business Week*, April 11, 2005.
- Hicks, Matt. (2005). Tags turning Web chaos into categories. *EWeek*, March 17, 2005. Retrieved June 5, 2005, from <http://www.eweek.com/article2/0,1759,1777004,00.asp>
- Hiler, John. (2005). *Google's war on hierarchy, and the death of hierarchical folders*. Retrieved July 15, 2005, from <http://www.microcontentnews.com/articles/deathofhierarchy.htm>
- Jacob, Elin K. (2003). Ontologies and the Semantic Web. *Bulletin of the American Society for Information Science and Technology*, 29(4), 19–22.
- Keller, Philipp. (2005). *Tags: database schemas*. Retrieved July 15, 2005, from <http://www.pui.ch/phred/archives/2005/04/tags-database-schemas.html>
- Koivunen, Marja-Riitta. (2001). *Annotea Project*. Retrieved July 16, 2005, from <http://www.w3.org/2001/Annotea/>
- Koivunen, Marja-Riitta; Swick, Ralph; and Prud'hommeaux, Eric. (2003). *Annotea shared bookmarks*. Paper presented at the KCAP03 workshop at Sanibel, FL. Retrieved July 16, 2005, from <http://www.w3.org/2001/Annotea/Papers/KCAP03/annoteabm.html>
- Leuf, Bo; and Cunningham, Ward. (2001). *The wiki way: quick collaboration on the Web*. Boston: Addison-Wesley.
- Library of Congress Cataloging Policy and Support Office. (2004). *Library of Congress subject headings*. 27th ed. Washington, DC: Library of Congress.
- Lund, Ben; Hammond, Tony; Flack, Martin; and Hannay, Timo. (2005). Social bookmarking tools (II): a case study – Connotea. *D-Lib Magazine*, 11(4). Retrieved June 12, 2005, from <http://www.dlib.org/dlib/april05/lund/04lund.html>
- Markoff, John. (2005). By and for the masses. *New York Times*, June 29, 2005, p. C1.
- Mathes, Adam. (2004). *Folksonomies – cooperative classification and communication through shared metadata*. Retrieved April 6, 2005, from <http://adammathes.com/academic/computer-mediated-communication/folksonomies.html>
- Merholz, Peter. (2000). [Scattered blast of thoughts on designing with information. Blast two: classification schemes]. Retrieved July 17, 2005, from <http://www.peterme.com/browsed/browsed040100.html>
- Mote, Nick. (2004). *The new school of ontologies*. Retrieved April 6, 2005, from <http://www.isi.edu/~mote/papers/Folksonomy.html>
- National Center for Supercomputing Applications. (2002). *NCSA Mosaic history*. Retrieved July 17, 2005, from <http://www.ncsa.uiuc.edu/Divisions/PublicAffairs/MosaicHistory/history.html>
- Parsia, Bijan. (2003). Semantic Web services. *Bulletin of the American Society for Information Science and Technology*, 29(4), 12–15.

- Quintarelli, Emanuele. (2005). *Folksonomies: power to the people*. Paper presented at the ISKO Italy–UniMIB meeting in Milan. Retrieved June 22, 2005, from <http://www.iskoi.org/doc/folksonomies.htm>
- Rosenfeld, Louis. (2005). *Folksonomies? How about metadata ecologies?* Retrieved April 6, 2005, from http://www.louisrosenfeld.com/home/bloug_archive/000330.html
- Rox, Torsten. (2005). Social bookmarks' [sic] charts version 3.0. Retrieved July 28, 2005, from http://www.irox.de/file_download/3
- Roush, Wade. (2005). Tagging is it: social bookmarking emerges as a wiser way to navigate the Web. *Technology review (Cambridge, Mass.)*, 108(6), 21.
- Schwartz, Candy. (2001). *Sorting out the Web: approaches to subject access*. Westport, CT: Libraries Unlimited.
- Shirky, Clay. (2003). *The Semantic Web, syllogism, and worldview*. Retrieved July 28, 2005, from http://www.shirky.com/writings/semantic_syllogism.html
- Smith, Gene. (2004). *Folksonomy: social classification*. Retrieved April 6, 2005, from http://atomiq.org/archives/2004/08/folksonomy_social_classification.html
- Sturtz, David N. (2004). *Communal categorization: the folksonomy*. Retrieved February 3, 2005, from <http://www.davidsturtz.com/drexel/622/communal-categorization-the-folksonomy.html>
- Taylor, Arlene. *The organization of information*. 2nd ed. Westport, CT: Libraries Unlimited.
- Weibel, S.; Kunze, J.; Lagoze, C.; and Wolf, M. *Dublin Core metadata for resource discovery*. Retrieved July 28, 2005, from <ftp://ftp.rfc-editor.org/in-notes/rfc2413.txt>
- Wright, Alex. (2004). *Folksonomy*. Retrieved April 6, 2005, from <http://www.agwright.com/blog/archives/000900.html>