

Why Are They Tagging, and Why Do We Want Them To?

by P. Jason Morrison

Folksonomies

Social tagging systems, which allow large numbers of users to classify items, are a fascinating subject. When I began planning for my master's thesis in the [Information Architecture and Knowledge Management Program](#) at Kent State University, I knew I wanted to study collaborative tagging and folksonomies. The next step, deciding on which systems to study and how to study them, was more difficult.

Just a year or two earlier the task of narrowing down the types of folksonomies to study would have been much easier. Now it seems as though nearly every website invites users to tag articles, photos, videos and merchandise with keywords. Folksonomies are no longer limited to cool Web2.0 startups working out of their garages to change the way we think about enterprise social wiki pet care. Amazon displays user tags on item pages, IBM sells products that employ social tagging and Microsoft's blogs have tag clouds.

In the course of choosing specific sites to study, I looked at a lot of different tagging systems and folksonomies. While browsing, searching and trying out as many tagging systems as I could get my hands on, I noticed some broader issues that I could not cover directly in my research but definitely merit discussion. Information architects and web developers will be asked more and more to incorporate these systems into websites, and there are a number of open questions about how the systems should work and how they are used.

On the surface, tagging systems and the resulting folksonomies seem almost magical. Is it really possible that the problem of organizing and

classifying can be solved as simply as allowing random users to contribute tags? This notion is made even more tempting by the fact that many content management systems now have the functionality built in, or it's just a matter of installing a plug-in.

I suspect, however, that a folksonomy is most likely to be successful when the goals of the website or information system intersect with the goals and motivations of users. So when considering how to add tagging functionality and how to build and use a folksonomy, we have to ask: What are folksonomies for, and why do users tag?

What are folksonomies for?

Folksonomies are generally used to organize information and support information retrieval (IR). The public face of many folksonomies is often a tag cloud, with tags usually listed alphabetically and weighted by popularity. Users might navigate the folksonomy by following tags they've used, related tags, popular tags or recent tags. In many cases folksonomies are also used to support search. A site with a huge database of photos like [Flickr](#), for example, might rely almost entirely on a folksonomy for search since there is no full text to fall back on.

Tags might have some other, more subtle uses as well. Even if users do not use the folksonomy to navigate the site, the presence of descriptive tags may suggest whether they are on the right track. In addition a site's information architects and user experience designers could make use of a folksonomy themselves. For example, tags could be data mined to provide additions to synonym rings or an insight into the kinds of topics users are thinking about at the moment.

It is also possible that a site might adopt tagging and create a folksonomy with no real interest in how well it supports IR whatsoever. Tagging could

P. Jason Morrison completed this study as a graduate student at Kent State University. He is also a senior analyst at AT&T. He can be reached by email at who<at>jasonmorrison.net

be seen as a small way to get users interacting with a site, to get them to return, to encourage them to sign up for a user account or just to give them a richer experience and a sense of participation. This is the same reason why many sites allow users to rate items, vote in polls and contribute comments.

When thinking about adding tagging to a site, the first question should be: What do we want to get out of this? Does the site need something to improve search results or a new navigational facet to better connect related pages? Is the goal to classify lots of multimedia objects with minimal cost or to get users to interact with the site a little more? The answer could be "all of the above." Of course if the cost of adding the functionality is negligible, the answer might even be "who knows, let's throw it at the wall and see what sticks."

Once the goals have been decided, decisions about the tagging interface can be made to support them. There are so many open questions at this point that are just begging for further research. Should the tagging system suggest tags to users? If so, should it suggest popular tags from other users or that user's most-used tags? Should the tagging system attempt to control vocabulary in any way or perhaps just apply spell-check? What should the search system search? When users bookmark a web page with [Furl](#), for example, they are able to contribute topics, keywords, comments, clippings and a rating, and [Furl](#) can also save a copy of the page.

In my research on the use of folksonomies to support search I found a great example of how relatively small details can impact the usefulness of the entire system. Time and time again I wondered why so many systems encourage or require single-word tags. On the popular tags page at [del.icio.us](#), tags like *howto* and *rubyonrails* illustrate that expressing even simple concepts often requires more than a single word. Users can work around this limitation by concatenating words together, but this likely limits the folksonomy's IR performance. While users browsing a tag cloud might recognize *howto* as denoting items that explain "how to" do something, users typing search queries will not have the same luck. [Del.icio.us](#) contains thousands of items tagged *information*, *architecture* and *informationarchitecture*. If a user searches for "information architecture," how is the search system to know

which of those three tags are relevant? If one of the goals of the folksonomy is to support search, my guess is it is worth the extra lines of code to support spaces or at least some substitute character.

Why are users tagging?

Social bookmarking systems like [del.icio.us](#) and [Furl](#) worked well for my study because users tagged websites in order to organize their bookmarks and find things later or to share them with others. On photo and video sites users might be motivated to tag for similar reasons but also because without relevant tags the items they are adding to collections may not be searchable at all. When I looked at news sites like [Slashdot](#), however, I began to wonder what utility users get from tagging. Tags on some sites were often inside jokes by regular users or comments on the quality of a story. On some sites when the headline for a story was in the form of a question the most popular tags were invariably "yes" and "no."

Although I did not test any news sites, my guess is that in cases like this the results are much weaker for search. Inside jokes might help some long-time users navigate, but it is hard to imagine the utility of a tag cloud filled with *yes*, *no* and *maybe*. Users must be getting something out of these systems, though, and it is important to realize that different users have different motivations for tagging.

Users tag things in order to find them again later. Users of social bookmarking sites like [del.icio.us](#) and [Furl](#) might use those systems to discover websites and share them with others, but the primary goal of bookmarking is tagging an item so that you can find it again. Tags can later be searched or used to organize a large collection into categories in tune with the user's own idiosyncratic mental model.

This scenario is one we usually have in mind when we talk about folksonomies. One user might tag a given photo with *sailboat*, while another chooses *schooner*, a third *Miami* and a fourth user might choose *peaceful*. Individually, the users are motivated to provide good keywords, and with enough users the folksonomy should become fairly robust.

Providing good keywords is not, however, the only motivation for users

to tag. My guess is that this particular case would likely result in folksonomies that are effective in information retrieval, so long as there are enough users participating.

Users tag things to get exposure and traffic. Content producers submitting items to a collection will tag their items so that users browsing through or searching the folksonomy will ultimately see their content. Bloggers, for example, tag stories with an eye toward showing up in [Technorati](#) searches.

In this case, tags are not completely different from the meta keywords that were once used by search engines in the early days of the web. Does this similarity necessarily mean tags submitted by authors will inevitably become filled with spam? Not necessarily. It certainly seems logical to prohibit abuse behaviors such as registering large numbers of fake user accounts to push up rankings or adding excessive numbers of keywords. My guess is that most folksonomies are not harmed when creators submit their items and tag them with a sane number of keywords, especially if the items would be virtually invisible otherwise.

They key difference between most folksonomies and the old, abused meta keywords system is that users can mitigate the impact of spam tags by providing their own, more relevant tags. Since most sites will have many users tagging for every content producer, the spam tags should be pushed down over time. Also, abusive tagging can lead to consequences for the spammer. This potential behavior might not be a concern at all if users are trusted, as in an intranet setting.

Users tag things as a way of voicing their opinions. Many social news aggregation sites like [Digg](#) and [Reddit](#) have voting or rating built directly into their interfaces, but when a system only includes tagging, users may provide tags that are a judgment of the content rather than a description.

Some sites like [Slashdot](#) even encourage this practice. Tags like *slownewsday* and *fud* regularly appear, meant by users to denote that a story is not very notable or as an attempt to sow fear, uncertainty and doubt, respectively.

It is a little hard to see how a folksonomy built from these tags would be effective in information retrieval. How often do people search for

“awesome” or “finally”? In this situation an information architect might recommend adding a rating system, since users are obviously interested in that kind of interaction.

Keep in mind that qualitative tagging is a perfectly valid activity, so long as it aligns with the goals of the site. A humor site might get much more value from a folksonomy generated from collaborative sarcasm than it would from a very accurate system of classification. These tags might even help IR in some ways, perhaps by enhancing the information scent of certain items. A search for “Firefox” might turn up thousands of results, but the results also tagged *cool* might seem like better paths to pursue than those tagged *lame*.

Users tag things incidentally as they perform other IR tasks. Some users may be tagging when they do not even know it. As part of an independent study project, I created a website called [Mealographer](#) that allows users to track the nutrition in their diets. In the course of evaluating the usability of the site, I found that the text of the USDA database of foods did not provide very good results for users’ full-text searches. I could, however, capture the text of the searches and then associate that text with the food item the user ultimately decided on. Users were tagging items implicitly when they searched and then found an item that satisfied their search.

This method can help in some situations – users searching for “salad” might see salad dressings, potato salads and similar items ranked highest. Once a user realizes that *salad* is too general a term, a search like “salad with lettuce and carrots” would allow common salad constituents to climb quickly through the ranks as later users searched for “salad” and then chose lettuce off the list.

I did find a number of possible drawbacks to this method. For one, if the item’s current description does not intersect with common search terms at all, it might never come up in search results and then never be tagged. For example, the entry for *cola* would never show up on searches for “Pepsi” or “Coke.” Over time the entry might rack up tags like “Coca Cola” or “RC Cola” but that still leaves a lot of Pepsi drinkers frustrated. The best technique to mitigate this issue I found was to cheat – to manually tag the items myself.

Users tag things to take advantage of functionality built on top of a folksonomy. This can be the case for either content producers or users. While some blog writers may employ Technorati tags primarily to gain exposure, Technorati also allows bloggers to tag their posts to create links to similar articles on other sites. Some plug-ins to popular blogging software like **WordPress** employ tags to generate lists of related articles from within the site. In recommendation systems like **StumbleUpon**, users may be motivated to rate or tag websites in order to get more interesting, targeted suggestions.

As folksonomies become more integrated into the functionality of websites this motivation may become more and more important. If content producers and users are motivated to tag in order to get related items, the resulting folksonomies could support IR quite well. It would be interesting to see if this would result in less varied tags – perhaps bloggers would realize that more posts are tagged with *IA* than *information architecture* or vice versa and change their tagging accordingly.

Users tag things to play a game or earn points. Although this application is not very common, it is a very interesting one. One great example is Luis von Ahn's **ESP** game. In the game, users are presented with images pulled from the web and asked to guess the same keywords as another user. If the two match, they get points and move on to the next image. Some sites already reward users for forum posts or other user-generated content, and it would be quite possible to do the same for users that tag prodigiously.

The resulting folksonomies may or may not support IR as well as the other scenarios already discussed. Users primarily trying to earn points might submit a lot of tags that are just good enough to not get them banned from the site. Community-driven sites could use additional functionality such as user ratings to reward quality tags and contributions to ameliorate this problem.

My experience in playing the ESP game was that my tagging behaviors changed as I tried to rack up more points in the game. I began using shorter words and simpler concepts to try to get more matches under the time limit, and I found myself not going with my first instinct and instead entering words I thought other players would use. Since I found that others were

typing “sky” if the sky was at all visible in a photo, I began tagging the same way, whether or not it seemed like a striking or important part of the image. Of course more study is needed, and these sorts of tags might be exactly what are needed for certain applications. When presented with a large and unorganized collection, games like this may just be the fastest and least expensive way to get a large amount of classification done.

Aligning user motivations with site goals

The discussion of uses for folksonomies and the list of motivations for tagging are far from complete, but the value of taking some time to match up site goals with user goals is already clear. It is also clear that information science researchers have a lot of work to do in the next few years to get an empirical grasp on the strengths and weaknesses of all the variations we have seen and the new ones being created in a dorm room as we speak. ■

Resources

Websites Mentioned in the Article

Information Architecture and Knowledge Management program at Kent State University – <http://iakm.kent.edu/>
 Slashdot tagging FAQ – <http://slashdot.org/faq/tags.shtml>
 del.icio.us – <http://del.icio.us/>
 Digg – www.digg.com
 ESPGame – www.espgame.org/
 Flickr – www.flickr.com/
 Furl – www.furl.net/
 Mealographer – www.mealographer.com/
 Reddit – <http://reddit.com/>
 Slashdo – <http://slashdot.org/>
 StumbleUpon – www.stumbleupon.com/
 Technorati – www.technorati.com/
 WordPress – <http://wordpress.org>