

Tags: What are They Good For?

Prentiss Riddle

riddle@io.com

<http://prentissriddle.com>

INF 385Q: Knowledge Management
School of Information
The University of Texas at Austin

5/12/2005

Introduction

Tags and folksonomies represent a novel approach to the problem of organizing information. These terms refer to the use of uncontrolled user-supplied textual labels ("tags") to categorize and identify resources in a nonhierarchical shared information space (a "folksonomy").

Although the concepts of uncontrolled keywords and user-supplied metadata are not new, the particular environments in which they are being used have generated intense interest; as of this writing Google shows 273,000 entries for the term "folksonomy," a neologism which did not exist a year ago. Yet at this time there are just two relatively mature and heavily used systems based on these ideas, the bookmarking system del.icio.us (<http://del.icio.us>) and the photo sharing system Flickr (<http://flickr.com>). That these two systems are the exemplars on which the enthusiasm for folksonomy rests is intriguing, because del.icio.us and Flickr are different enough to show that not all folksonomies work in the same way and yet they are similar enough not to necessarily bound the space in which tagging might be useful.

This paper will briefly state some of the strengths and weaknesses of tagging that have been demonstrated by existing systems, but its focus will be to speculate about the larger, as yet unanswered questions regarding other domains in which tagging may or may not prove useful.

What we know so far

Millions of words into the folksonomy frenzy, what do we know? Existing systems have demonstrated that tagging is good for a handful of things.

Tagging is useful for **personal recall**, or finding again what you have seen before. This is the main purpose of the bookmarking system del.icio.us and its many less widely used imitators (see Irox (2004) for a comparison of several).

Tagging supports **social effects**. One frequently noted example is the common practice in Flickr of agreeing on a tag and convening an informal group of users who create new photos expressly to use it. The canonical instance of this practice is the tag "squared circle" (<http://flickr.com/photos/tags/squaredcircle/>).

Tagging promotes **serendipity**, or the pleasant and sometimes useful discovery of the unexpected. Prowlers of libraries and used bookstores know that the collocation of books on a shelf by topic creates interesting juxtapositions of items which are sufficiently different that the information seeker would not have thought to explicitly search for them. Tagging heightens this effect by allowing items to be "collocated" on any number of serious or whimsical criteria, including the identities of the users who tagged them, and by increasing the factor of chance through the imprecision of tags.

Tagging is good for **novelty**, or what del.icio.us creator Joshua Schachter calls "interestingness" and defines as "the first derivative of popularity" (Weinberger, 2005, p. 20). Schachter's statement implies that ideas or resources which are already widely known become uninteresting through familiarity and that items which are gaining in attention at a given time are those most likely to provoke our interest. Tags provide a convenient way to segment and measure the rise and fall in popularity both of individual resources (how many times they have been tagged) and the categories they fall into (the tags themselves). This focus on the new is presumably why the default presentation of items in del.icio.us and Flickr is in reverse chronological order and it dovetails nicely with the common support of RSS feeds in tagging systems.

We also know that tagging (at least as it stands now) is *not* good for several other sorts of information discovery and retrieval.

Tagging does not perform well at measures of **precision** and **recall** (in a different sense from "personal recall" above). In information retrieval systems, *precision* is a measure of the elimination of false positives, that is, the portion of retrieved documents which are relevant. *Recall* is a measure of the elimination of false negatives, that is, the portion of relevant documents in the system which are retrieved (Korfhage, 1997, p.194). Because users can apply many different tags to a single concept or a single tag to many different concepts, the results for a given query are likely to be both noisy (low in precision) and incomplete (low in recall).

Tagging is not good for **ontologies** in the sense of precisely defined relationships among concepts. The most familiar sort of ontology in the domain of information discovery and retrieval is a hierarchical system for classification by subject with a controlled vocabulary and a thesaurus defining broader and narrower terms. The uncontrolled vocabulary of a tagging system is essentially the polar opposite of a system like the Library of Congress Subject Headings or the Dewey Decimal System (Mathes, 2004). Similarly, free tags do not provide the structure necessary to capture and enforce compliance with other kinds of ontologies, from faceted classification systems to the schemas being constructed to support the Semantic Web (for example those at <http://www.schemaweb.info/>). It is true that users of del.icio.us or Flickr can create their own conventions to represent hierarchies or other relationships, sometimes expressed by internally segmenting a tag with punctuation, but since there is no formal mechanism for recording or validating those conventions and since the tags are recorded in a common space with overwhelmingly noncompliant tags, such effects are local.

What we don't know yet

There is a longer list of things about tags and folksonomies which we do not yet know. I will discuss these questions of their possible utility in terms of the domains to which they may be

applicable, the social and organizational contexts in which they may work, and the degree of structure which may be imposed on them.

Domains

Besides bookmarks and photos, what other domains might tagging be good for?

People. If folksonomies are this year's great fad, it could be said that last year's was social network systems. Opinions vary as to the success of commercial social network services such as Friendster and Orkut, but an area of continued interest is distributed social network systems expressed in XHTML (XFN, n.d.) or XML schemas such as FOAF (FOAF, n.d.) and the Trust Ontology (Golbeck, 2004). While these systems use a restricted vocabulary of labels to identify social relationships, it is easy to imagine their extension to include uncontrolled tags (in fact, given their distributed nature and the lack of a centralized system to impose adherence to the standards, it is hard to imagine that they would not be widely extended).

Another novel use of tags to apply to people is the online dating service Consumating (<http://consumating.com>), in which users tag themselves and each other through an interface resembling that of Flickr. The Consumating model could easily be applied to the tagging of people for other purposes, such as employment, staff directories or sales contacts.

But tagging people raises the question of whether tags' infamous ambiguity would become problematic when it is applied to their social relationships. As a simple example, at Consumating the tag "tall" can mean that a would-be dater *is* tall or that he or she wants a tall *partner*. The directional if not hierarchical nature of many social relationships might put pressure on such systems to move back toward more structured ontologies. And then there is the question of people's likely sensitivity to the tags others apply to them. While an unflattering tag in del.icio.us might irritate the author of a web page to which it is applied, the same tag applied

to a person might be perceived as a direct attack, creating resentment of the system hosting the tag.

Products. Large retail systems like Amazon and eBay already permit extensive commentary and description of products by their customers and affiliated vendors. They could easily welcome tags as an additional way to describe and locate merchandise. However, a vendor with a closer branding and ownership involvement with its product line might feel differently. For example, imagine the corporate reaction if a tagging system at Disney.com were used by the community of fans who write erotic fiction about Disney characters!

Places. There has been interest for some time now in user-contributed metadata associated with geographic locations, such as the GeoURL project (<http://geourls.org>), but as originally conceived the idea suffered badly from scalability issues (Riddle, 2003). Augmenting geolocation information with tags might help address these scale problems. A tagging interface for ubiquitous GPS-enabled camera phones could mean that we will soon "annotate the planet" as Job Udell memorably says (2005).

Music. Users of iTunes and its associated library of shared metadata, Gracenote, often complain about the limited number of genre categories the library supports. iTunes itself accepts other user-defined genres, which can pass from one iTunes user to another when embedded in MP3 files, but there is no way to assign a song to more than one genre or to explore songs by user-defined labels outside the confines of one's own MP3 collection. Support for tagging could be a welcome addition to iTunes or similar music software. The aggregation of tags in services like the Apple Music Store or Gracenote could be equally welcome, although privacy concerns stemming from the contentious issue of filesharing might require strict anonymization of any

uploaded metadata that would limit the social features common among other folksonomy systems.

Filesystems. Apple Spotlight supports a keyword feature which resembles tagging, but because keywords are generally accessible only to a single user they do not constitute a folksonomy. The Spotlight model does, however, raise the possibility of cooperative tagging of shared files on a LAN or multiuser system.

Tasks. The extension of del.icio.us-style tagging to documents or pages in intranets or organizational websites is an obvious idea to pursue. Less obvious is the idea of using tagging to identify tasks or workflows. Where a task corresponds to a page with a fixed URL, the problem is straightforward; but for tasks that require authentication, temporary session IDs or URLs specific to each user, an "enterprise del.icio.us" would not be sufficient. Nevertheless, designers of content management systems may find it worthwhile to support tags which map across session- or user-specific boundaries, particularly if tagging proves effective in the static parts of an intranet.

Social and organizational context

So far tagging has proven itself in the context of free, Internet-wide systems of 60,000 to 300,000 users (del.icio.us and Flickr, respectively) and one to four million documents (Weinberger, 2005). Questions remain about the other contexts in which folksonomies might work.

Scale. Does tagging perform better at some scales, and under what circumstances? For environments in which searches need not be exhaustive and serendipity is part of the fun, there may not be an upper bound to the size of a tagging community. On the other hand, a tagging system intended to serve a specific purpose such as commerce or navigation within an enterprise

may be more sensitive to increases in noise that might accompany increases in scale. In such a case, a possible solution might be to filter infrequently used tags out of displays, to reduce the visible tags for a particular item or the items for a particular tag to a manageable number.

Conversely, the benefit of aggregating individuals' tags into a shared system may be lost if the community is too small. It remains to be seen whether there is there a lower bound at which social effects diminish and a folksonomy might as well be a single-user system.

One concern about folksonomies is that they might simply be an early-adopters' club. We do not know whether the current interest in tagging is a fad that overstates its benefits and will fade with time. What's more, early adopters' ease at creating social norms and informal protocols may have substantial effects on a system's friendliness and signal-to-noise ratio which disappear once they are joined by large numbers of less sophisticated and less community-minded users, a phenomenon colloquially known as the "the AOL effect" or "the death of Usenet." We do not know whether Flickr would still be a friendly place with three million users instead of 300,000.

Organizations. If tagging depends at present in part on the "solipsism" or immediate self-interest of taggers (Veen et al., cited in Lawley, 2005), can that translate to a business or organization? One might be less or more likely to put one's bookmarks on display before a closed group of peers than the world at large. As in other instances of knowledge management, how these questions play out in practice depends on subtle matters of organizational culture and the economics of information which are not easy to manage or predict (Davenport, 2000).

Structure

If we find that the unconstrained nature of tagging itself introduces constraints on the applications for which it is suited, one avenue to consider would be to adapt folksonomies to be slightly more structured.

Normalization. A number of people have suggested that tagging might benefit by an interface which helps people re-use their tags consistently or use them consistently with other taggers (for example Ducker, 2005). However, Joshua Schachter has objected strongly to such proposals because they would violate the reliance on "intuition" which he feels is essential to the mnemonic function of del.icio.us (Weinberger, 2005).

Ease of use. It is commonly observed that tagging succeeds where other distributed metadata systems fail because of the low cost of tagging in terms of the user's time and effort. It is possible that better user interfaces might put normalized or modestly structured metadata within the same ease-of-use threshold, if not for entering metadata then perhaps for retrieval. Recent work on Rich Internet Applications (RIAs) suggests that complicated tasks involving the narrowing of decision trees can be made simpler through interfaces which are not limited to the click-and-load web model (Schleicher et al., reported in Riddle, 2005). Tools for easier navigation of hierarchical or faceted classification systems such as Endeca (Weinberger, 2005) might offer models for investigation even though they would not be directly applicable to the less structured systems of folksonomies.

Information retrieval methods. Like any other form of metadata, tags are potential fodder for information retrieval systems. So far the contribution of IR to folksonomies is mostly a matter of speculation, although Flickr does include a "related tags" feature which appears to be based on a technique more sophisticated than raw terms counts. We do not yet know whether tags can be usefully clustered or disambiguated through aggregation, nor whether tags would be

a useful addition to the data considered by ordinary web search engines. If the tags, taggers and URIs are not sufficient information with which to apply IR methods, then the content of the tagged items and the taggers' social networks are also available. All of those possibilities make the topic complex enough that we may not see clear answers soon.

Directions for research

The many questions raised above naturally call for further research to find answers, but the available methods may be as thorny as the problems they are trying to solve.

The Cranfield method traditionally used in IR research of testing search methods against small, predetermined data sets seems unlikely to be adaptable to a process which is less algorithm-centered than user-centered and social. Traditional usability research might be more applicable, but again the social nature of folksonomies means that models based on studying users in isolation may tell us about only a small part of the system. Ethnography and other non-quantitative methods will no doubt be useful, but they may be limited by applying to existing systems whereas many of the questions we want to answer are about hypothetical ones.

In his survey of tagging systems, David Weinberger (2005, p. 31) concludes his own list of unanswered questions by saying:

There is a simple solution, however, to all of these issues: Create the tags and experiment. Tags are becoming a new layer of infrastructure. They will enable yet another round of creativity as we figure out, collectively, what variety of things we can do with this metadata.

His point is correct as far as it goes, but that "Field of Dreams" approach is not an entirely satisfying answer to investors deciding what particular kind of tagging system to put

their money into, nor to the creator of a production system with a job to do, nor to a user trying to decide which of several competing systems will be worth his or her time spent on tagging.

One partial answer comes from the open-source community. Possibly the construction of feature-rich and hopefully interoperable playgrounds for early adopters, like the abundant but chaotic state of browser plugins, would sustain the interest of enough users to inform decisions about more focused systems for production environments.

References

- Davenport, T.H., & Prusak, L. (2000). *Working knowledge: How organizations manage what they know*. Boston: Harvard Business School Press.
- Doctorow, C. (2005). *Folksonomy, or how I learned to stop worrying and love the mess: impressionistic transcript*. Summary of a panel discussion at the O'Reilly Emerging Technology Conference, San Diego, CA. Retrieved April 10, 2005 from <http://craphound.com/etech2005-folksonomy.txt>
- Drucker, A. (2005, January 4). *A modest proposal*. Posted to <http://slashdot.org/comments.pl?sid=134812&cid=11253049>
- The FOAF Project (n.d.). Retrieved May 1, 2005 from <http://foaf-project.org/>
- Golbeck, J. (2004). *The FOAF Project*. Retrieved May 1, 2005 from <http://trust.mindswap.org/trustOnt.shtml>
- Irox (2005, April 22). *Social bookmarks charts version 2.2a*. Retrieved May 1, 2005 from <http://www.irox.de/stat-pdf/socialbookmarks.pdf>
- Korfhage, R. (1997). *Information storage and retrieval*. New York: Wiley.
- Lawley, L. (2005, March 13). *SXSW: Leveraging solipsism*. Summary of a panel discussion at South By Southwest Interactive, Austin, TX. Posted to http://www.corante.com/many/archives/2005/03/13/sxsw_leveraging_solipsism.php
- Mathes, A. (2004). *Folksonomies - cooperative classification and communication through shared metadata*. Retrieved April 10, 2005 from <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>

Riddle, P. (2003, January 21). Austin blogging density approaches infinity, hell sighted 0 miles SE. Retrieved April 10, 2005 from

<http://www.aprendizetodo.com/toys/?item=20030121>

Riddle, P. (2005, March 7). RIAs and the death of the page. Retrieved May 1, 2005 from

<http://prentissriddle.com/blog/index.php?p=19>

Udell, J. (2004, August 20). Collaborative knowledge gardening. Retrieved May 1, 2005 from

http://www.infoworld.com/article/04/08/20/34OPstrategic_1.html

Udell, J. (2005). A Google Maps walking tour of Keene, NH. Retrieved April 10, 2005 from

http://weblog.infoworld.com/udell/gems/gmap2_flash.html

Weinberger, D. (2005, February). Taxonomies to tags: From trees to piles of leaves. *Release 1.0*, 23, no. 2, 1-33.

XFN: XHTML Friends Network (n.d.). Retrieved May 1, 2005 from <http://gmpg.org/xfn/>