

Folksonomy Folktales

During a recent project, I had occasion to review the latest writings about folksonomies in general, and comparing folksonomies and taxonomies in particular. I've been a bit leery of a lot of the claims for folksonomies, but wanted to see what, if any, new ideas there might be. What I found was, on one hand, some interesting experiments in combining taxonomies and folksonomies, and on the other hand, a whole lot of very enthusiastic writings about why folksonomies are better than taxonomies. What I didn't find was a very good discussion of the relative strengths and weaknesses of the two. . Instead rather there was article after article repeating the same myths, folktales, and misconceptions.

So, what I'd like to do now is take a look at some of these myths, folktales, and misconceptions and try to dig a bit deeper and ultimately see if we can't come up with a more realistic view of both folksonomies and taxonomies.

Before we get into specific myths and misconceptions, I'd like to first take a look at a fundamental flaw in the vast majority of articles on folksonomies and taxonomies which is the almost universal use of the Dewey Decimal system (or Library of Congress Subject Headings) as the example taxonomy. Using the Dewey Decimal system as your example taxonomy when trying to discuss the pro's and con's of taxonomy and folksonomy says to me that either you have no understanding of taxonomy creation and use in today's world or you are just so set on showing the superiority of folksonomies that you have set up a rather silly strawman to then gleefully knock the stuffing out of.

It's like if the question you are exploring is the relative merits of jet skiis versus boats and you pick the latest model of Jet ski and talk about how great it is, how nimble, how quick, how cheap, and they swarm all over the place and are really fun. And now for boats, you pick as your example – the Titanic. It's really big and cumbersome and made of brittle steel held together with bad rivets and it costs too much and it's too hard to build and it's slow and hard to steer and runs into icebergs and kills lots of people. OK, so jet skiis are better than boats – or jet skiis have all these great characteristics but yeah, if you can afford it the Titantic has more comfortable beds.

But wait, sail boats are boats too – and they are much smaller, cheaper, easier to build, and lots of fun. And you know what – there are lots of taxonomies that are smaller than the Dewey decimal system, easier to construct and use, less rigid, have built in revision procedures and user input capabilities, and generally don't suffer from all those “characteristics” of taxonomies that folksonomy advocates love to list.

So think about this -- if the Dewey Decimal system is the only example of a taxonomy you can think of, maybe you shouldn't be writing about taxonomy and folksonomy. Or at least do some more research into the kind of smaller, more flexible, more responsive taxonomies that are being developed. There is more to taxonomy than the Dewey Decimal System.

Another overall impression I got from a review of new and old articles on folksonomies is that most of the articles are guilty of a rather massive overhype showing a great deal of enthusiasm but not so much careful thought. For example let's take a look at the opening of an often cited article:

"The Hive Mind: Folksonomies and User-Based Tagging
by Ellyssa Kroski

There is a revolution happening on the Internet that is alive and building momentum with each passing tag. With the advent of social software and Web 2.0, we usher in a new era of Internet order. One in which the user has the power to effect their own online experience, and contribute to others'. Today, users are adding metadata and using tags to organize their own digital collections, categorize the content of others and build bottom-up classification systems. The wisdom of crowds, the hive mind, and the collective intelligence are doing what heretofore only expert catalogers, information architects and website authors have done. They are categorizing and organizing the Internet and determining the user experience, and it's working. No longer do the experts have the monopoly on this domain; in this new age users have been empowered to determine their own cataloging needs. Metadata is now in the realm of the Everyman. "

I have to admit that my first reaction was, "Oh no, not another revolution! Didn't we just have one last year and a couple the year before?" Maybe I'm getting jaded, but I really think that we should be a bit more careful to not cheapen a very good word. The printing press was a revolution. The Industrial Revolution was a revolution. The Internet is an ongoing revolution. But Folksonomies? I think not.

Note: In the following discussion we will distinguish folksonomy from user-generated tags. The focus will be on folksonomy because that is what is mostly being written about and it has what many writers believe is the essential characteristics of social feedback. One reason for this is that user-generated tags are not tied so completely to the idea of some sort of order emerging from people simply seeing how others are tagging.

Aside from the revolutionary fervor, this quote also exemplifies several of the standard general folktales about folksonomies.

Folktale One: Folksonomies are examples of the wisdom of crowds.

Actually folksonomies are the exact opposite of the wisdom of crowds. If you read James Surowiecki's book, *The Wisdom of Crowds*, the key characteristic of situations in which you get a wisdom of crowds effect is that no one can be aware of what anyone else is doing. The reason that a crowd of amateurs can guess the weight of a bull better than an expert is that every guess is completely independent of what the crowd is doing. If you publish the guesses as they are being made, what you get is not the wisdom of crowds, but the madness of crowds – the bandwagon effect. Which means that folksonomies that publish tag clouds of popular tags do not exhibit a wisdom of crowds effect.

Of course, you're free to use the phrase, Wisdom of Crowds, to mean something else – perhaps that throwing a lot of people at a problem, regardless of how you set it up, will inevitably lead to a good outcome, but in that case, I would suggest that you read some history starting with Tulip Mania and going through to the latest real estate boom and bust, and for good measure, throw in the book, The Cult of the Amateur by Andrew Keen for a counter view.

Folktale Two: Folksonomies are building bottom-up classification systems.

First of all, folksonomies are not a classification system, they are an unordered, flat set of keywords that are ranked by popularity. Ranking words by their popularity can tell you a great deal about how groups of people are thinking and that information can be extremely useful, but it does not tell you much of anything about the relationships between words or concepts. In other words, there is no “onomy” in folksonomy.

Folktale Three: Folksonomies are Working

This one is somewhat in the eye of the beholder, but it seems to me that the claims for success for folksonomies are vastly overstated. Yes, there was a flurry of activity when Delicious and Flickr first came out and that a number of similar sites sprang up on the internet. However, the growth of new sites seems to have hit a plateau and a closer examination reveals that the number of people actually adding tags remains pretty small. For more on this question, take a look at the sections on the limits of folksonomies.

Also, there are two parts to “working”. First, folksonomies as tagging discussed above. But there is a second sense, which is about using folksonomies for actually finding information and here the answer is clearer – they are not all that powerful a mechanism for finding information. Browsing for like-minded people and using folksonomies for serendipitous browsing is fun and occasionally useful, but it represents a very small percentage of search behavior.

Folktale Four: Metadata works best when it is free – in the realm of Everyman

This seems to me to be basically pure ideology and expresses more of a cultural and political philosophy than an actual claim (See the section Why the Fuss later in this article). I haven't seen any evidence for this claim and the experience of information architects and librarians asking people to tag documents in an enterprise environment strongly suggests the opposite. If you have to choose between metadata tags generated by users and authors without any guidance from a taxonomy versus those created by IA's and librarians from a controlled vocabulary or taxonomy, the answers were pretty clear that the latter produced much better results.

On the other hand, it was always a struggle to get people to add metadata at all and if we again restrict our attention to the actual act of tagging, then yes, folksonomies seem to have an advantage in that it is easier to get some people to just think of a tag off the top of their head than to select a value from a complex taxonomy. But there are two caveats. First, as we see from the

number of people tagging, folksonomies don't give you the kind of coverage we were looking for outside of general social bookmarking sites, especially within the firewall of enterprises – that is, getting everyone to tag. So folksonomies work better for getting some people to tag, but not for getting everyone to tag. This is another example of the difference between the Internet where getting anyone to tag is a plus and the enterprise/intranet where having only some documents tagged is not an answer. I'd like to see more research on just how many people are tagging at the different sites and who is doing the tagging.

Also, while I will grant that it is probably easier to think of a tag than select from a complex taxonomy, I'm not so sure that it is easier than selecting from a simple taxonomy. I'd like to see a lot more study on that one, something that goes beyond the visceral dislike of taxonomies by a few well known authors and the use of our old strawman, the Dewey Decimal System.

Anti-Taxonomy Myths

Let's shift from our general folktales to more specific claims and myths about taxonomies and folksonomies. To start let's look at another frequently cited article on folksonomies that is a good source for the standard myths about the drawbacks of taxonomies. This same list appears in other articles, so I hope I'm not guilty of erecting a strawman to knock down. Many of these "drawbacks" are also found in one form or another in one of the most often cited articles in this area, Clay Shirky's [Ontology is Overrated: Categories, Links, and Tags](#)

[Folksonomies: power to the people](#)

Emanuele Quintarelli

Drawbacks of hierarchical schemes (taxonomies)

- 1 - Items do not always fit exactly inside one and only one category.
- 2 – Hierarchies are rigid, conservative, and centralized. In a word, inflexible.
- 3 – Hierarchical classifications are influenced by the cataloguer's view of the world, and, as a consequence, are affected by subjectivity and cultural bias.
- 4 – Rigid hierarchical classification schemes cannot easily keep up with an increasing and evolving corpus of items.
- 5 – Hierarchical classifications are costly, complex systems requiring expert cataloguers to guess the users' way of thinking and vocabulary (mind reading)
- 6 – Hierarchies require predictions on the future to be stable over time (fortune telling)
- 7 – Hierarchies tend to establish only one consistent authoritative structured vision. This implies a loss of precision, erases difference of expression, and does not take into account the variety of user needs and views.
- 8 – Hierarchies need expert or trained users to be applied consistently

This is a well organized and often repeated list of taxonomy drawbacks. There is just one problem: Pretty much every single one of these is wrong, misleading or overstated, or a known issue for which taxonomists have worked out methods to overcome over years of practice.

Let's take a closer look at this list of taxonomy characteristics:

“1 - Items do not always fit exactly inside one and only one category.” Yes, this is true, but taxonomists and ontologists have been dealing with this one for a long time and there are lots of ways to handle various situations, including polyhierarchy and other standard techniques.

“2 – Hierarchies are rigid, conservative, and centralized. In a word, inflexible.” No they aren’t. Flat out wrong. Some are, many are not. Virtually every taxonomy that I or a number of friendly competitors have developed are designed to be flexible (offering alternatives), progressive (built in maintenance plans to reflect change in users and/or corpus), and are hybrid models that include both a central team and constant input from users. Also, a number of studies including ones I’ve done show that folksonomies at sites like Del.icio.us are surprisingly conservative, with little change in most popular tags and much fewer new terms than you might expect.

“3 – Hierarchical classifications are influenced by the cataloguer’s view of the world, and, as a consequence, are affected by subjectivity and cultural bias.” True, but overstated. Hierarchical classifications are influenced by the cataloguer’s view of the world, and, if well done, by the view of the world of a variety of other cataloguers, current theory of good classification, and dozens to 100’s of world views of potential users. Plus, we have developed a variety of methods to combat the influence of personal and cultural bias.

Finally, just how are Folksonomies free of bias? Communities by their very nature harbor prejudices and bias. Del-icio.us, for example, is dominated by one culture – high tech computer people and that bias produces things like the tag, “Design” to point almost exclusively to bookmarks about software design.

“4 – Rigid hierarchical classification schemes cannot easily keep up with an increasing and evolving corpus of items.” False or overstated – depending heavily on the framing words “rigid” and “easily”. How about flexible hierarchical systems with built-in procedures for adding new terms? True, these procedures require a level of effort but what is “easy” varies depending on the context. They are certainly a lot easier than most folksonomy advocates might imagine.

“5 – Hierarchical classifications are costly, complex systems requiring expert cataloguers to guess the users’ way of thinking and vocabulary (mind reading),” Overstated. First, note again the emotional frame words – “guess” and “mind reading”. Sorry, but that’s bogus. It is not guessing or mind reading when you follow well researched and tested methods for obtaining input into the user’s way of thinking and their vocabulary. However, it is true that hierarchical classifications require expert cataloguers but they needn’t be complex and as far as being costly, while it is true that taxonomies cost more to develop than a folksonomy in traditional ways of measuring cost, there are two caveats. First, in corporate environments within the context of enterprise search, the cost of developing a taxonomy is a very minor cost. Second, the low cost of a folksonomy is based on a myth: that user’s time is worthless and therefore free. Try adding up all the time that users spend tagging bookmarks and add that to the cost.

“6– Hierarchies require predictions on the future to be stable over time (fortune telling).” False and misleading. Again, note the derogatory “fortune telling”. Taxonomies only need to reflect

the present (and often the past) and include a mechanism to handle change and novelty. This is something that all good taxonomists do, but which seems not to have occurred to too many folksonomy enthusiasts.

“7 – Hierarchies tend to establish only one consistent authoritative structured vision. This implies a loss of precision, erases difference of expression, and does not take into account the variety of user needs and views.” Mostly false. Hierarchies do tend to establish a single consistent authoritative structured vision, but they can and usually do include variations. The supposed drawbacks are in the second sentence and they are essentially wrong and seem to be based on a really odd view of the role of taxonomists.

We’ve already discussed how good taxonomists always do an enormous amount of research into user needs and views. In fact, given the tyranny of the majority effect you often see on folksonomy sites, I’d say they probably do a better job of handling the needs and views of the users who are not part of the majority community. Next, “erases difference of expression” – no they don’t, they preserve those differences in a variety of ways, including representing minority views directly in specially designed parts of the taxonomy to the simple mechanism of variant terms, all of which can be exposed.

“8 – Hierarchies need expert or trained users to be applied consistently.” Overstated. Yes, there are known issues with consistency of tagging with a taxonomy, those problems are massively multiplied when using no taxonomy or, the same thing, a folksonomy. Also, consistency is a sliding scale – the more complex a taxonomy, the more difficulty with consistency and as we have seen taxonomies are not all Dewey Decimal System level of complexity. Finally, there are software options that can reduce the need for expert users.

So in summary, the 8 anti-taxonomy myths are:

- 1 – true, but with established and effective counter measures
- 2 – False
- 3 - Overstated with established and effective counter measures
- 4 – False
- 5 – Overstated
- 6 – False
- 7 – Mostly false
- 8 – Overstated

Benefits of Folksonomies

So if the majority of anti-taxonomy myths are false or overstated, what about the other side of the story, the perceived benefits of folksonomies in general and in comparison with taxonomies?

- 1 – Folksonomies are easy to use.

The trouble with this one is, as indicated above, it depends on what you are looking at. Here we have to sharply distinguish between the act of tagging and the use of tags to find information.

It certainly seems on the face of it that just picking a word to apply to a bookmark or document is a cognitively easier task than reading through a complex taxonomy and deciding which term to apply. However, I believe that the ease of tagging with top of the head terms is over-stated. And the corollary, that the cognitive task of picking a term from a taxonomy is easier than claimed. And, of course, selecting from a simple taxonomy (something that folksonomy advocates don't seem to believe in or know about), is even simpler. And when we move from general internet sites to either targeted vertical internet sites or enterprise sites with some control over the content, the balance shifts even more, especially when we add in categorization software that can make suggestions from within the taxonomy. In that case, the cognitive task of agreeing with the suggestion or not is much easier than trying to think up a term.

And finally, there is the issue of quality, that is, is it easier to generate good and useful tags off the top of your head than to select from a taxonomy? And while that gets us into difficult waters with deciding what a good tag is, it doesn't seem to me that you can avoid the question.

Which brings us to the second sense of easy to use – are folksonomy tags easier to use to find information? Having no internal structure and no relationships between terms, makes it much more difficult to use tags to find information in a variety of ways. First of all, let's take an example of where I'm looking for information about a topic such as neuroscience. I did a study of LibraryThing, a social bookmarking site for librarians, and discovered that with the lack of even such simple variants as plurals, that I would have had to click on about 15 different terms to get a fairly complete coverage of the just the high level overview of the field of neuroscience.

Some of the issues were plurals (neuroscience and neurosciences), and related terms like (cognitive neuroscience, cognitive psychology, etc.). And even more fundamental was the relationship between the general term neuroscience, and the sub-topics within neuroscience. In a folksonomy there is no relationship and so each sub-topic is independent and has to be selected separately. The most general terms in a folksonomy are simply that – general, in other words, they don't include specific sub-topics. And so the answer might be that that's OK, the tag is just used to refer to bookmarks that are only general neuroscience. Unfortunately, that quickly falls apart when we take a closer look, in part, because people who don't know a subject very well will always choose the most general term whether the bookmark is for general neuroscience or on synaptic strengths in learning.

2 - Tags more accurately reflect population's conceptual model

I see two immediate problems with this. First, it's not clear that most users have a coherent conceptual model, certainly nothing like the coherence of a taxonomy. In one sense of the word, yes, everyone has a conceptual model – or rather they have multiple conceptual models some of which are quite deep and detailed and structurally as rich as any taxonomy, but some of which are small, skimpy, and fragmented. For example, early research on experts found that they typically had a rich collection of about 50,000 elements (combination of facts and concepts and relationships) and that these elements are very well structured in way that allow experts to chunk

large aggregates that they can bring to bear on a problem as a single entity. These structures are probably richer and more useful than any formal taxonomy.

A problem, however, arises when we start talking about non-experts. And even experts are non-experts in some field. And here the picture is radically different. Non-experts tend to have much fewer facts and relationships and the structures are very fragmented and which fragment gets applied to a given situation is often largely determined by outside influences. For example, someone might tag a bookmark with the tag X the first time they see it, but a virtually identical bookmark might be tagged with Y because in the meantime, they have been exposed to other ideas that influence their categorization. (The influence of one categorization schema on learning another schema is called “Intertwinedness” in category theory.) And this is particularly a problem in that there is no practical mechanism for normalizing tags over time.

So the second point is that reflecting an incoherent conceptual model is not really the best way to put a metadata tag front end on a content collection. This is true even for someone using their earlier tags to find their own information and even more true for other people trying to use that tag to find new information. This is not to suggest that people are stupid and librarians should rule, but rather that categorizing and tagging is a specific skill that most people are not trained in.

A related problem is that experts and non-experts categorize differently. To explain, let’s look at an intriguing idea from category theory which is called a natural level or basic level category. An example of a basic level category is the word, “dog” which in common taxonomies appears below mammal and above a number of types of dogs like Golden Retriever or Boxer.

Basic level categories are intermediate levels within a hierarchy that have a number of important characteristics. First, they are categories that children learning a language tend to learn first and they are categories that tend to be used more often. Second, they are categories whose members have a particularly powerful combination of expressiveness and distinctiveness. Distinctiveness refers to how strongly members can be distinguished from other members on the same level. So in our example, dogs are very different from cats – much more different from each other than say a Golden Retriever and a Golden Lab.

Now, aside from a fascinating (for some of us) digression into category theory, what does this have to do with folksonomies and reflecting user conceptual models? Well, the answer is that experts tend to have a preferred level of categorization that is lower or more specific than a non-expert. So, for example, a dog show judge would not use the word “dog” for a picture of their favorite toy poodle, they would use the lower level, toy poodle. On the other hand, someone who knows very little about a subject tends to tag a level higher than average. They might choose the word “philosophy” for a web site that discussed a number of philosophical issues about epistemology while a more experienced reader might choose the middle tag of “epistemology” and an expert might choose a lower level tag like mind-body problem.

If all you want to do with folksonomies is use them for community building in which experts would likely self-select and the differences in categorization wouldn’t matter as much, then that is fine. It is an interesting way for like-minded and like-experienced people to find each other. The problem is when you start comparing folksonomies and taxonomies in terms of usefulness in

finding information. They are not comparable in any significant way. Discovering people with the same interests and same level of expertise is great, but it is a very small part of information behaviors like searching for content.

3 – lack of a central authority

Being out from under the thumb of those dictatorial librarians seems to be something that folksonomy advocates very devoutly wish for. My guess is that it has more to do with ideology than any actual tagging and/or searching activities and is strongly associated with a technology culture that mistrusts academia and especially anything that smacks of the humanities.

However, regardless of why folksonomy advocates hate the central authority of librarians so much, there are good reasons for not taking this claim at face value. Let's start with the idea that is often cited that folksonomies "allow a decentralized collaborative view to emerge". This sounds refreshing and liberating but in reality, this decentralized view tends to be surprisingly rigid and conservative. If you track the top 50 or so tags at sites like Delicious or others, they have actually changed little over the last 3 years. And in studies that I've done, this conservativeness holds true at lower levels also, once you get a minimum number of tags and tagged objects.

More importantly than the surprising conservativeness of popular sites like Delicious, is the idea that these popular sites do have a central authority, and it is the most oppressive and most dangerous type of central authority there is, the authority of the majority. There is a reason that we don't see true democracies in politics any more – they are the most dictatorial form of government there is. Against the will of the people, there is no recourse, no way of insuring the rights of the minority.

To move from politics to tagging, think of this example, the tag "design" has been in the top 25 most popular tags for the last three years. Now the word design has a number of meanings and there are a number of different kinds of design, but an examination of the web sites tagged with design on Delicious reveals an almost exclusively small set of meanings having to do with computer or software design. This is an indication of the nature of the Delicious community, which is dominated by computer geeks as opposed to word geeks or artistic geeks or scientific geeks. Now suppose that you, like what looks like about .01% of Delicious users, want to tag your favorite interior design site with the word "design" because either the system doesn't handle noun phrases or because you just find it easier to use single word tags. And now suppose you want to use the tag "design" to find other sites that deal with interior design or other people that are interested in interior design. Good luck! If you click on the word "design" in your basic tag cloud, you will get approximately 1 million hits 99.9% of which have nothing to do with interior design. Happy wading through page after page of computer design links.

Now, there are clearly ways to get around this, but they all entail having to compensate for the overwhelming impact of the majority (the collaboratively emergent majority). This makes it harder to tag because no one but you is worried about the rights of the minorities. And even if you devise a workaround, what are the odds that the other interior designers will think of the

same workaround? And as the number of self-selected computer geeks swamp Delicious with their version of what “design” or other terms mean, this effect will grow even more pronounced.

It seems to me that having a system in which there was a central group of authorities or librarians that you, as a minority can appeal to might work a bit better than letting the collaboratively emergent dictatorial majority unconsciously ride roughshod over the minorities.

Another type of authority that emerges on social sites is that of well-connected individuals and fashion leaders. I would like to see an in-depth study comparing this type of authority with a designed central authority but my fear is reflected in the W.B. Yeats line – “The best lack all conviction, while the worst are full of passionate intensity.” As long as there are mechanisms for detecting these passionate worst (from the simple ability of people to detect bullshit to editorial staff on Wikipedia) then the worst of this problem can be at least alleviated. However, given the madness of crowds structure of social networking sites, I’m worried that this will get worse or at least settle down to a mediocre middle dominating fashion like they do on TV.

4 – Quality of tags is 90% of taxonomy

Now this one can probably just be dismissed as wishful thinking or mindless ideologically inspired nonsense. First, there have been to my knowledge no studies comparing the quality of tags on folksonomy sites with the quality of tags selected from a taxonomy that could be used for the basis of such a numerical claim. Early studies of tagging inside the firewall suggest that 90% is wildly exaggerated. I could with as much justification say that they are about 50% or 10% as good.

5 – Better than nothing at all – often accompanied by argument of expense and difficulty of getting people to use hierarchy

When pressed, such as an early exchange between Clay Shirky and Lou Rosenfeld, it often seems that the last defense that folksonomy advocates fall back on when the issue of quality of tags is raised, is well, they are better than no tags at all. This is often accompanied by the argument that since taxonomies are so expensive and so difficult for people to use, at least folksonomies will get people tagging and even poor tags are better than no tags. Aside from the overstatement of the expense and difficulty of using taxonomies to tag, there are grounds to doubt that any tags are better than none. I remember a study that demonstrated that just asking people to tag as they wanted to actually made documents harder to find.

It is important to keep in mind that tagging is not an end in itself. It is an activity that is meant to lead to using those tags to find information. And if bad tags, make it harder to find information, then having more tags is not a good thing. At the very least this claim needs a lot more study before we can claim that some tags are necessarily better than none.

I also find it striking that we often go from folksonomy being a revolution to “well some tags are better than nothing at all”. That doesn’t sound all that revolutionary to me – especially when you move outside the realm of general internet bookmarking sites.

From my perspective, I would like to end this benefits section with what I think could be one of the more significant benefits from folksonomies and the discussion about them, which is that just getting people excited about metadata is a huge plus. The real trick will be to maintain that level of excitement while enhancing the quality and usefulness of user generated tags.

Revolution? Limits of Folksonomies

The problem I have with folksonomy advocates like Clay Shirky and others is that they wildly overstate pretty much everything about how revolutionary folksonomies are, how difficult and limiting taxonomies are, and worst of all, the comparison of taxonomies and folksonomies.

For something that is touted as revolutionary, folksonomies have a number of major limitations which severely restrict their impact both in terms of the kinds of uses that can be made of them and the types of web sites and environments that they are useful on.

We've already covered a number of limitations on the use of folksonomy in specific sections above, but let's summarize them here. One limit which is one that all information strategies have to deal with, is the issue of scale. When you only have a few thousand or few million web sites in total and the number of links to any single term is in the hundreds or thousands then almost any information strategy will work. The problem comes when you start getting a million links to popular tags and large scale web sites like Delicious are already reaching that level for more and more of their tags. Back when the internet was first really getting started regular search engines worked pretty well, but once the number of pages got large enough, the quality of the results got worse and worse. I expect the same thing is already happening with large general sites like Delicious.

And there is another side of the scale issue, which is that on vertical portal sites, the amount of content tends to be small enough that user generated tags might not be as easily overwhelmed with millions of hits, but on the other hand, the target for searching for content on these sites tends to be much more specific, with the focus on finding one or a few particular documents, not a general web site. This means that the scale issue is here also.

Another limit on the applicability of folksonomies is that virtually all the success stories are from Internet sites, but when the approach is applied to intranets, the value proposition seems to go way down. Actually, within the enterprise, user-generated tags are being explored and are being found to generate a lot of value – but only when adapted to an environment that tends to include the central authority of information professionals working with users. (See the section, Knowledge Architecture for Enterprise 2.0). Folksonomies with their emphasis on having no central authority are much less valuable in an enterprise environment.

There is one more scale limit on folksonomies, that can be seen in a study done that compared LibraryThing, a social tagging site for librarians, and Amazon. What they found was that both offered social tagging but that even though Amazon had 1000's of times the number of users, they generated far less tags. The suggested reason was that people only tagged if they had about 200 or more items to keep track of. Less than that and there was no reason to tag for personal benefit and without a personal benefit, very few people bothered to tag much of anything.

Finally, another limit to folksonomies has to do with the well documented poor quality of tags on social networking sites – issues like plurals, misspellings, synonyms, idiosyncratic and extremely personal only tags like toread, funny, box47, or on the table. These kinds of tags severely limit the social usefulness of tags.

However, there is another quality issue which is more significant in my mind and that is there is really no mechanism for improving the quality of tags that people will do. First of all, there is no evidence that seeing how other people tag leads someone to become a better tagger. There are some suggested ways that seeing other people's tags might make your tags better, but I would argue that you are just as likely to learn bad habits as good tags (again good in the sense of finding info). And, in addition, there is no evidence that anyone actually goes back and retags. This would seem to go against the whole notion of the ease of tagging that is credited with getting people to tag. Indeed, many sites don't allow retagging.

Lastly, it is very useful to remember that, as enthusiasts often forget, the majority of people writing about folksonomies tend to be extremists. Not just in terms of their enthusiasm for folksonomies, but also in terms of their overall behavior. A journalist or thought leader is someone who by nature or by job is happy to spend hours every day playing and exploring and surfing and connecting – because that is what they love and/or that is what they are being paid to do. But most of the world is not like them – something that many, so many writers forget. Yes, the new generation spends more time connecting than the last, but no one, outside of the enthusiasts, spends as much time connecting, and thinking about connecting, and thinking and talking about connecting and tagging as the enthusiasts.

Of course, to be fair, I should mention that librarians have been just as guilty of assuming that everyone was like them and the result was a whole lot of advanced search applications designed by librarians that no one but a librarian could or would use.

What are folksonomies good for?

Does all of the above anti-folksonomy rant mean that I think that folksonomies are completely useless? No, not at all. For certain limited areas and applications they are a useful and even exciting but hardly revolutionary addition.

Within the limits discussed above, folksonomies can be a very useful mechanism for their core strength – discovery of other people who have similar interests. For discovering and exploring other people with similar interests and sharing the bookmarks that they have added and thereby greatly expanding your resources in those areas of interest is very cool and an exciting way to learn and connect to people.

An added benefit that we talked about earlier, is if you take into account the propensity of experts to tag at lower or more specific levels of tags, then experts can create a self-selected community automatically. People who tag something with “philosophy” are ignored and other experts who tag with things like “categorical imperative” know automatically that others who tag with this level of specificity are likely experts too.

Within the limits discussed above, folksonomies can function as a mechanism for discovering and creating communities – the heart of social networking. And, while I still would not use the word revolutionary, clearly social networking is very important and exciting and is changing the dynamics in a number of areas from marketing to politics and this is a real benefit.

Knowledge Architecture for Enterprise 2.0

The real value of folksonomies or at least user-generated tags can best be realized, however, by locating them in a more sophisticated and developed infrastructure. This approach is one that I have seen or been part of within the enterprise and it follows a very common pattern which is that new ideas are often developed on the Internet (and often overhyped) but then see their more complete and valuable development within the enterprise. The key idea behind this approach is to develop a framework and a facilitation capability within which user generated tags and folksonomies can be supported and utilized in ways that go way beyond tag clouds for serendipitous browsing for potentially interesting communities. This approach is what I call knowledge architecture and it is an approach that has a great deal of potential for enhancing KM in the enterprise in general and for enterprise 2.0 initiatives in particular.

A full description and discussion of knowledge architecture is well beyond the scope of this article, but some of the main elements as they relate to folksonomies and enterprise 2.0 include:

Knowledge Organization elements:

- Taxonomies, metadata of all sorts, but particularly faceted metadata and best bets
- Categorization rules and name catalogs or entity dictionaries
- Maps of formal and informal communities within the organization and their associated knowledge and information behaviors

Software:

- Text analytics, enterprise content management, enterprise search, KM platforms

People Resources:

- A central KA group including librarians and knowledge architects – who develop, refine, and maintain the infrastructure elements and facilitate their application

Just to get a flavor and inkling of what a knowledge architecture approach to folksonomies and enterprise 2.0 might involve, let's look three areas: Development of knowledge organization and community maps, Application in ECM and Search, and Refinement of knowledge organization

Research Resources:

First, in the area of development of such knowledge architecture elements as taxonomies and community maps, folksonomies can be used as resources for social and cognitive research into the interests and language of users and communities. We already do this to some extent with search log analysis but user generated tags are a higher order cognitive task and represent a level of categorization higher than simple search terms. This makes them a much more powerful resource, especially if they are used in conjunction with search log analysis.

Another area where user generated tags can be a great source is for developing faceted navigation applications. Research on popular book marking sites reveals that the vast majority of tags fit very well with a faceted classification system. For example, the article by Louise Spiteri talks about the overwhelming use of thing tags. My own analysis of Flickr, shows that close to 80% of the tags fall into one of a few facets – thing, place, event, and people. So if you are in the process of developing your facets, user generated tags are great sources for specific values and an analysis of a folksonomy created by your users is a great source for designing which facets are particularly important to your users.

Search and ECM Applications:

Combining folksonomies with other elements is also a way to get more direct benefits than using them for research. The way to do that is to develop a hybrid approach that combines folksonomies with simple taxonomies. The taxonomy need not be the Library of Congress classification – it could be a 2, 3 or 4 level taxonomy at most.

The basic idea is that the simple taxonomy guides and offers suggestions to the user who can just take the taxonomy term, take the term and add another tag for more specific level tagging or just add their own tag independent of the taxonomy.

This type of hybrid normally has an editor function so these user-generated tags can be analyzed and either treated as suggested new terms for the taxonomy or simply sub-categories. A hybrid approach, while not satisfying the revolutionary ideologues, nevertheless, offers real value and is one way to extend the application areas of folksonomies beyond sites like Delicious.

Another possibility within the enterprise is to develop a 3 or 4 level taxonomy that is supported by content management software that can have categorization rules associated with each taxonomy node so it can offer full suggestions based on its auto-categorization functionality. This makes the taxonomy piece even more powerful but can also incorporate user suggestions and additions.

This hybrid approach is much more powerful for finding documents and internal web sites, but can it also support social networking? If there is a software front end that functions like a tag cloud with tags ranked by popularity but also referencing the taxonomy, it should enable social networking that is at least as good as a straight folksonomy based site. I'd like to see a lot more research on this, however, on both enterprise and internet sites.

Finally, another area that a knowledge architecture approach works well is simply creating the infrastructure to support a range of KM or Enterprise 2.0 initiatives. Despite some over simplified and over hyped rhetoric, 2.0 initiatives don't just happen. It is not enough to buy some blogging and wiki software and make it available to your employees and sit back and watch them transform your organization. We've already seen how taxonomies and folksonomies can support and enhance each other, and when it comes to things like wiki's there is a similar need for knowledge organization and editorial roles to get their full benefit.

For example, all but the most simple wiki applications need an initial intellectual framework or categorization schema to avoid chaotic structures that limit the usefulness of a wiki. They also

need ongoing editorial functions to monitor the wiki, resolve disputes, and ensure adherence to a level of content quality. For even the granddaddy of wiki and 2.0 applications, Wikipedia utilizes thousands of editors for these basic editorial functions.

A well-developed knowledge infrastructure and central KA group is also the way to get the maximum benefit from blogs and other KM 2.0 initiatives. This can range from utilizing a taxonomy-folksonomy hybrid as part of the front end to a constantly expanding internal and external blogosphere to using the KA group to facilitate the use of taxonomies in publishing to using the knowledge organization to integrate external blog content into your internally generated content.

This just gives a high level view of how knowledge architecture can be used to expand the application and usefulness of folksonomies. Of course, while I prefer the approach and term, knowledge architecture, there are a number of terms and approaches to develop a taxonomy-folksonomy hybrid. For more, see my upcoming book on Knowledge Architecture.

Deeper Themes – why the fuss?

The last question I want to look at is – why the fuss? What led to this over-hyping of folksonomies? A full cultural-psychological analysis is way beyond the scope of this article which has already gotten too long, but let me just suggest a couple of themes that it seems to me had something to do with it.

First, there is simply the fact that social networking is really very cool and important and is having a lot of impact in a lot of areas. So it is very easy to go from “It is very cool to be able to connect to other people all over the world” to “The mechanism that enables that connection must be very powerful too.”. IOW, if social networking is revolutionary then folksonomies must be revolutionary too. What quickly follow from this is article after article elevating (and overstating) the power of folksonomies.

A second factor that is related to the first is simply an IT/technology focus on tools. Let’s face it, we’ve been using the Internet and other means for social networking for some time. What is different is that the current set of tools, particularly blogs, wikis, and social tagging sites make that social networking easier, faster, and available to more people. All of which is great and opening up social networking and tagging to larger masses of people does have an impact (both good and bad as scale effects become stronger), but it’s only a revolution rather than just the next evolutionary step if you think that tools are the really fundamental component.

A third factor can be seen when we look at how folksonomies are being elevated. We’ve already examined the specific arguments, but here I want to look at the underlying philosophy/culture behind those arguments. People in IT and other technical areas have very different knowledge cultures than people in so called soft sciences like the humanities or “management science”. I am currently working on a book on knowledge architecture in which I go into these different knowledge cultures in more detail, but in general, there are different types of concepts dominant in IT/technical cultures and more importantly for this discussion, different criteria of “truth”. To

oversimplify, in IT/technical cultures, something is true if it works – the algorithm produces a correct answer or solves a problem. Who creates that algorithm is irrelevant to the correctness of the answer – it can be a 13 year old kid playing in their room or a 20 year veteran with myriad honors and a huge reputation. In the humanities or management, it matters much more who says something because there is no universally accepted way of objectively evaluating the correctness of most statements.

One outcome of these differences is that there is a great deal of mistrust and antagonism toward the importance of reputation and experience in the soft sciences by IT/technical people which often comes out as a strong anti-authoritarian, anti-expert attitude. It is actually not so much anti-expert as anti-non-technical expert. In the area of folksonomies that is often accompanied by a suspicion of culture concepts like semantics. To see how this plays out, let's take a look at the [Ontology is Overrated – Categories, Links, and Tags](#) article by Clay Shirky in which he discusses the drawbacks of the early Yahoo browse the Internet categorization schema.

“But Yahoo decided to privilege one way of organizing links over all others, because they wanted to make assertions about what is “real”.”

I know some of the people involved in the early Yahoo categorization schema and I can tell you this is way off the mark. Yahoo cataloguers weren't trying to model reality or impose some Nazi view of reality on poor defenseless users. They were trying to help users find content better – and make some money while doing it. The point was that no one could really find anything on the Internet – it was one of the first inflationary periods and Yahoo gave them a way to find stuff faster and easier – not a big conspiracy theory attempt to impose rigid thought control on the world of free taggers. Who of course, at the time didn't tag at all.

And another quote:

“Yahoo is saying, ‘We understand better than you how the world is organized, because we are trained professionals. ...but the effect was to override the users' sense of where things ought to be, and to insist on the Yahoo view instead.’”

Here the anti-authoritarianism comes through loud and clear – or rather certain types of experts. But again, this is kind of silly. Let's be realistic – it's not as if you had all these alternative ways floating around the Internet of developing categorization schemas to organize all content on the Internet. I know that I didn't meet more than 2 or 3 people whose stated goals were categorizing all Internet content in their spare time.

And as in the notion that they were attempting to model reality, the idea that Yahoo cataloguers were actually claiming that “they understand better” is way off the mark. First, of all, we need to distinguish between claims. I would imagine that most people who have had training in categorization would never make such a sweeping claim, What they would claim, with a great deal of authority, is that they know more about categorization and developing general purpose categorization schemas than someone with no training or experience. However, good cataloguers would never make the second part of the claim – that they better understand how the world is organized – because they have enough experience to understand that there is no single

correct way to organize the world. The best you can come up with is something that has certain good formal properties and reflects a pragmatic compromise that according to their user testing, works pretty well.

Yes, this compromise will not accurately reflect how every single user would have done the job and so every user will have some aspect of their categorization overridden by the Yahoo (and other) categorization schema. And this is also true for every folksonomy as well insofar as it incorporates any kind of relationship between tags. Folksonomies only avoid this problem by not offering any sort of categorical relationship whatsoever.

A third factor in the “revolutionizing” of folksonomies might have arisen from a combination of one aspect of our IT/technical culture which is the mechanism of “releases” turned into a metaphor – 2.0 in this case, and the standard marketing hype of everything is revolutionary (from a new soap to a new more efficient way to publish content). To counteract this particular (bad) marriage of concepts, I recommend reading two books, *Metaphors We Live By* – George Lakoff and Mark Johnson, and *Death Sentences: How Cliches, Weasel Words, and Management-Speak are Strangling Public Language* by Don Watson.

Conclusion

First and foremost is that Folksonomy should never have been compared with taxonomies, thesauri, or ontologies. The comparison is essentially bogus in that folksonomies are not a classification system at all. Popularity is not a semantic structure – to repeat, there is no onomy in Folksonomy. I give the creator of the term, Thomas Vander Wal, full marks for creativity and marketing acumen, but failing marks for accuracy.

Folksonomies do not organize information, they are aggregates of 1,000’s or million’s of individual acts of cataloging ranked by popularity. Yes, there are conceptual relationships between articles and web sites that are expressed when two or more are tagged with the same tag and these relationships can grow in very complex ways revealing a great deal about how people think and how some ideas are can be related. But it’s not clear how the overall set of conceptual relationships constitutes an organization of knowledge. This becomes particularly clear as the number of tags and sites multiply and the complexity of the tag and community relationships grows exponentially. Gee, haven’t we seen this process before?

On the other hand, it is possible to browse through tags and citations and pick up a number of ideas of how other people have tagged a particular set of articles/web sites and thus be exposed to a variety of connections between concepts. This can be fun, it can be rewarding, but it is ultimately rather limited in its impact.

Perhaps the most fundamental limit of folksonomies for future “revolutionary” impacts is that there is really no mechanism for improving the quality of the individual tags that is consistent with their most powerful characteristic, ease of use. I just don’t see flocks of users going back and re-tagging or even necessarily getting better at tagging simply by being exposed to other people’s tags, in fact many sites don’t allow it. The enthusiasts might very well learn to use

folksonomies to improve their own tags, but I really wonder how widespread that behavior will be.

And so, the second conclusion is that folksonomies are interesting but not really very revolutionary. User tagging has been around for a long time and the relationship between free keyword tagging and classification schema or taxonomies has been and will continue to be a dialogue in which hybrid systems will almost always be the best answer.

So, in my opinion, the future of folksonomies lies not with the revolutionary rhetoric which actually detracts from the real value of folksonomies but with a range of new and interesting ways to combine user-generated tags and taxonomies. To get an idea of what is going on in the hybrid world, take a look at some of the articles cited in the bibliography.

True, these new hybrid systems of taxonomy/classification schema and folksonomy won't satisfy the anti-authoritarian strain of thought that runs through folksonomy literature, but I will predict that they will ultimately deliver much more value to many more people in many more environments than pure folksonomy applications ever will.

Bibliography

[Ontology is Overrated – Categories, Links, and Tags](#) by Clay Shirky

The standard starting point for folksonomy discussion

[Folksonomies: power to the people](#) by Emanuele Quintarelli

A very good summary of the standard claims for folksonomies and the drawbacks of taxonomies

[Ontology of Folksonomy: A Mash-up of Apples and Oranges](#) by Tom Gruber

A good discussion of social tagging, ontologies (distinguished from Clay Shirky's use to mean hierarchical classification schemas), and a variety of semantic issues.

[Collaborative thesaurus tagging the Wikipedia way](#) (v2;2006-04-27) –[[Wikimetrics]] research papers, volume 1, issue 1. by Jakob Voss [Fascinating discussion of combining thesaurus categories and user tagging.

[The Structure of Collaborative Tagging Systems](#) by Scott Golder and Bernardo Huberman

A very good article analyzing tagging behavior and various cognitive and semantic aspects of tagging and taxonomy, including some of the surprising patterns over time within sites like Del.icio.us, CiteULike, and Connotea

[Creating the academic library folksonomy: Put social tagging to work at your institution](#) by Xan

Arch

A call for combining library and folksonomy

[Cataloging and You: Measuring the Efficacy of a Folksonomy for Subject Analysis](#) by Tiffany Smith

An analysis of folksonomies at LibraryThing, including various quality issues

[Beneath the Metadata: Some Philosophical Problems with Folksonomy](#) by Elaine Peterson

An interesting discussion of some of the underlying philosophy behind folksonomies (strong relativism)

[Structure and form of folksonomy tags: The road to the public library catalogue](#) by Louise Spiteri

A good analysis of folksonomy tagging including the overwhelming use of thing tags, and that for libraries to use folksonomies they have to design ways to overcome the problems with bad tagging.

[Taxonomy Directed Folksonomies: Integrating user tagging and controlled vocabularies for Australian education networks](#) by Sarah Hayman and Nick Lothian

A good overview analysis of folksonomies and user tagging, including ideas for combining folksonomy and taxonomy.