

Advances in classification research, Vol. 17: Proceedings of the 17th ASIS&T SIG/CR Classification Research Workshop (Austin, TX, November 4, 2006), ed. Jonathan Furner and Joseph T. Tennis.

AN EXAMINATION OF AUTHORITY IN SOCIAL CLASSIFICATION SYSTEMS

Melanie Feinberg

University of Washington, USA

1. Introduction

Champions of social classification praise its flexible and collaborative nature, in contrast to the rigidity and authoritarianism that they see in traditional classificative structures (such as Kroski, 2005; Shirky, 2005c; Merholz, 2004). In the view of these writers, social classification applications such as the photo storage Web site Flickr and the Web bookmarks manager del.icio.us are both *democratic*, incorporating the participation of all Web users, and *emergent*, changing rapidly in response to new content. On the other hand, traditional methods for organizing information, particularly those that involve hierarchy, are seen as exclusive, because they may not represent all users' viewpoints, and imprecise, because they cannot be easily adapted for the rapid pace of content development engendered by Web publishing.

Two claims appear to underlie these descriptions of social classification. One, that the goal of classification is to identify and locate items based on a personal sense of appropriate categorization, and two, that, if enough other users index (or tag) items according to their own personal ideas of appropriate categorization, then all possibilities will be represented, and both searching and browsing will be facilitated.

This paper will evaluate these claims, particularly in regards to the role and nature of authority in organizational schemes, and the intersection of authority with an organizational scheme's purpose. I consider these issues for three services often associated with social classification systems:

- Indexing of personal collections.
- Sharing of indexed personal collections.
- Merging of personal collections into a group-indexed aggregate collection.

The bookmarks manager del.icio.us is the primary example of a social classification system used throughout this paper.

2. Indexing of Personal Collections

Indexing of personal collections, while a feature of social classification systems, is a personal information management problem, and not in itself social. However, personal indexing provides a strong motivation for using a social classification system, and these systems are predicated on the assumption that if enough people index resources using terms that are useful to them, that the aggregation of these terms will be useful to others. Golder and Huberman assert, for example, “that users bookmark primarily for their own good, not for the collective good, but [this activity] may nevertheless constitute a useful public good” (Golder and Huberman, 2005).

Certainly it seems given that the user of a service like del.icio.us will assign index terms that are personally familiar and that accurately describe some aspect of the content, based on the indexer’s particular perspective. However, the usefulness of assigned index terms for personal information management may diminish as the user’s collection increases in size and as the purpose for which the resource was originally collected varies. The general purpose associated with personal information management on the Web is keeping and refinding identified resources (or, according to Bruce, Jones, and Dumais, 2004, keeping found things found). However, the general category of refinding a resource may encompass a variety of specific purposes, such as the following for Web bookmarks:

- Keeping frequently used resources handy (for example, bus schedules for routes often taken).
- Grouping resources needed for a particular project (such as planning a vacation).
- Identifying potentially interesting but non-critical resources (an article on climate change, for example).
- Locating occasionally used resources that are difficult to recall (such as the mailing list archive for the SIG-CR).

The often dynamic nature of specific purposes within the general realm of refinding can affect the utility of assigned index terms as the context changes (for example, as a potentially interesting issue evolves into an active project, and then to a completed one). In terms of personal indexing,

for example, the term “toread” is only useful until the identified document is actually read, and the term “wishlist” only until the longed-for item is obtained. At that point, for maximum utility regarding the individual user’s personal collection, index terms should reflect the current status of the resource (or, if the resource is no longer of interest, it should be removed from the collection).

The appropriateness of index terms may also change as the size of a collection increases and as a user becomes more knowledgeable about a subject and requires more distinctions in organizing those resources. For example, one might initially index resources with the general term “metadata” and then proceed to distinguish between items about the Semantic Web, the Dublin Core standard, social classification systems, and so on. The user has two options here: either assign *only* more specific terms to new resources, or assign more specific terms *in addition* to the previously used general term (“metadata”). If the index terms for older resources are not edited (and in del.icio.us, re-indexing for multiple terms is relatively onerous), then in the first situation, the older and newer resources on the same topic are effectively separated (for example, older resources on the Semantic Web remain indexed by “metadata”; newer ones are indexed by “semanticweb”). In the second situation, older resources, as well as more general resources for which more specific terms are not identified, become more difficult to locate as the number of resources indexed by “metadata” increases. On its own, the usefulness of the initial index term “metadata” markedly decreases. While this may be true for any postcoordinate indexing system, and one can envision workarounds for it, any solution would require either re-indexing, significant forethought, or both. Part of the appeal of tagging, however, is the low overhead associated with it (Shirky, 2005b).

Although the utility of initial index terms may degrade for personal information management purposes, the effect on users of the aggregated system is more complex. In the case of terms that may expire for the initial indexer (tovisit, wishlist, etc.), another user may still be interested in places that somebody once wanted to visit or items that somebody once wanted to buy. But the system then becomes as much archeological as current. In addition, while the difficulty of browsing resources indexed with single, general terms increases more rapidly when considering the aggregate collection, the power associated with searching combinations of terms may increase. Then, however, the system becomes search-focused, and the aspect of “serendipitous browsing” often associated with social classification decreases. Quintarelli writes, for example, that “starting from a tag (proposed by the user or others) it is possible to discover all items, from all users, that match that same tag simply with a click” (Quintarelli, 2005). However, in

del.icio.us, popular, general terms such as “java” and “design” are assigned to hundreds of thousands of disparate resources, making it almost impossible to browse productively without additional refinement. While Guy and Tonkin assert that “the consensus among those in the folksonomy community is that extra tags are always better,” it seems that there are definite tradeoffs in such an approach. (Guy and Tonkin, 2006). While it is more likely that an item I might want will be indexed with a term that I might use, it is also less likely that I can locate that item in the extremely large set of results, and many of the results may not be interesting to me.

3. Sharing of Indexed Personal Collections

Users of systems such as del.icio.us can view the sets of bookmarks and associated index terms collected by other users. While this sharing may be regarded as “social,” the organization and composition of resources is still a matter of an individual user’s personal authority. The experience is rather like looking through someone’s bookshelves.

However, while one uses personal index terms to refind collected items, one uses someone else’s terms to discover new items of potential interest. The behavior here is similar to the concept of chaining as described by Ellis, where information seekers follow chains of citations to locate potentially useful documents (Ellis, 1989). The difference is that instead of using a seed document to identify additional documents or authors for investigation, one uses a seed document to find people who indexed the document similarly, and then to explore the other documents that person collected and organized.

However, especially given the inefficiencies associated with browsing tags, even for an individual user’s collection, this can be a time-consuming process. Just as in citation chaining, when seekers learn to identify trusted authors or publications and to use those distinctions to prioritize resources to consider further, users of social classification systems may be reluctant to invest the time in exploring too many collections from unknown users. As Liz Lawley states, “I don’t want to know what the average eight-year-old calls an image. I want to know what my friends and colleagues call an image. Or a link. Or a photo” (Lawley, 2005b). If other determinants of authority are to complement personal acquaintance as a motivation in exploring others’ collections, then perhaps systems such as del.icio.us should consider means through which users can gain reputations for their collecting and indexing, or means through which users can proclaim relevant expertise or

other information that might help information seekers determine which personal collections are most likely to be of interest.

4. Merging of Personal Collections into a Group-Indexed Aggregate Collection

The result of merging personally indexed collections is where much of the interest in social classification seems to fall, as this is where the “democratic” and “emergent” qualities become most salient.

The “hive mind,” also referred to as “the wisdom of crowds” (from a book by the financial journalist James Surowiecki) or “collective intelligence,” is often valorized by social classification enthusiasts. The hive mind forms the basis for claims of “democracy” (for example, in Kroski, 2005). The essence of the hive mind idea is that the combined knowledge of a group of people will be more accurate than the knowledge of any individual, even an expert individual. While the editor of a controlled vocabulary may miss a term that a particular user might associate with a concept, a wide user base constantly adding and applying terms will be more likely to include it. In addition, this broad user base will add new terms to the system quickly, bypassing the lag associated with formal vocabulary development (this is the emergent quality).

While some, such as Golder and Huberman (2005), describe social classification systems as “collaborative,” the cluster of index terms that is generated for a particular resource is really an aggregate of individual decisions, not a cohesive collaboration. In contrast to some other applications of collective intelligence, individual knowledge is not actively synthesized in social classification systems. For example, proponents of the hive mind often cite examples such as the ability to accurately estimate weights and measures by averaging a group of individual guesses (this is the initial anecdote in Surowiecki’s book). Another example is the use of wikis to create collaborative documents, where users can edit others’ work (theoretically, to correct and improve it) and discuss changes. Social classification systems provide no equivalent means of synthesis.

While Golder and Huberman do demonstrate that the index terms assigned to a particular resource in del.icio.us converge on a relatively stable set of core tags, this stable aggregation does not represent “consensus,” as they describe it (Golder and Huberman, 2005). Consensus, as a form of decision making for groups, requires negotiated agreement and compromise amongst those involved, as well as a means to define the group and its values, and there is no means to

achieve this in current social classification systems. While social classification systems may be democratic in terms of being open to all participants, this characterization does not extend to any sense of the community coming together to determine how resources should be indexed, even by voting (although recommendations of “popular” tags at the point of initial indexing, as del.icio.us does, may represent this to an extent). If any political characterization of authority is accurate, social classification is more like libertarianism, where everyone’s whims are allowed to flourish.

The libertarian aspect of social classification enables its “long tail.” (Those acquainted with information retrieval research will recognize the long tail as the right side of the Zipf distribution.) While the long tail is a familiar concept in statistics, it was recently popularized by Chris Anderson of *Wired* magazine to describe an economic strategy for Internet businesses, with the idea that the market for a large number of not-so-popular items (the long tail of the demand curve) is the same or greater as the market for a small number of popular items (Anderson, 2004). Amazon derives a good deal of its sales from the long tail, for example: although only a few people may buy an unpopular book, there are many such transactions in aggregate. For social classification, the long tail comprises those idiosyncratic terms that few users apply to a resource. The long tail makes it more likely that a term will be used in the system and that it will be assigned to an item in the same way that a particular searcher would use it. This may be true; however, as there is no way to determine the semantics behind the way a term is used, an information seeker might not be able to take advantage of this property. For example, I could not find resources associated with social classification in del.icio.us by using the term “indexing,” because the term “indexing” has been used in too many different situations, and there were too many false drops. (Appropriate resources may or may not be tagged with “indexing”; the point is that I looked through pages of results and gave up before I found any.)

Of equal concern, the libertarian characteristics of social classification, in addition to enabling the long tail, create a hive mind that, instead of being without bias, as described by Kroski (2005), let all biases thrive simultaneously and yet without clear articulation. For example, the “popular” resources indexed with “feminism” in del.icio.us include an article on how “feminized” education shortchanges boys and one on how “feminism has destroyed the real men.” As a feminist, this is not what I expect, or want, to see when I use this term, but I can’t limit my search to exclude those semantics that I don’t care about or don’t agree with. The “popular” resources indexed with the “taxonomy” tag include many documents about folksonomies (social classification systems) some of which are critical of taxonomies, and some of which mention taxonomies only

tangentially. It may certainly be worthwhile to have a wide range of opinions on various subjects appear through an organizational scheme; however, it places a burden on the information seeker to determine if each item explains how to use taxonomies, or is critical of taxonomies, or is really about a subject related to taxonomies, and so on.

Clay Shirky, a vociferous advocate for social classification, posits that the use of different vocabularies itself will mark certain communities in the midst of the hive (Shirky, 2005c). He claims that people who use the tag “cinema” would prefer to be separate from people who use the tag “movies,” for example. However, as noted by Peter Merholz and others, there is little to support such an assumption (Merholz, 2005). The fact that some people use the term “movies” and some use “cinema” *may* mean little more than that, or it *may* indicate a completely different perspective such as that which Shirky describes. More likely, it may indicate both at once, as a variety of different uses of the terms are indiscriminately mixed together. Current implementations of social classification lack mechanisms to make the expression of a community perspective explicit (for example, a means by which users could precisely define the semantics of terms such as “movies” and “cinema”). This missing aspect is analogous to Patrick Wilson’s idea of the skeptical librarian, who, while not presuming to judge between authorities, is nonetheless able to identify and explicate them, in order for the user to make a better judgment (Wilson, 1983). A similar idea is expressed in Tuominen, Talja, and Savolainen’s (2003) proposal for multiperspective digital libraries. According to Tuominen, Talja, and Savolainen, collections of resources should display relationships between perspectives on a subject; the authors present an example that shows different resources as agreeing with, contesting the validity of, or representing an alternate approach to the subject of chiropractic.

However, while an explication of vocabulary semantics is not the same as controlling vocabulary, it’s unclear how this could be done without exerting some sort of control. Even to say “for some people, the term ‘classification’ appears to indicate X. For some people, it appears to indicate Y,” is to delimit the concept space and apply a certain judgment. Even if anyone could contribute to such explications, how would disputes be resolved? What if the way that an index term was assigned was argued to be demonstrably wrong, for example? On what basis would such a determination be made, and what would be done about it? One could claim that to even attempt such “clarification” would go against the very nature of social classification, that the role of skeptical librarian requires a certain authority that goes against its ideal of libertarian freedom. Guy and Tonkin (2006) warn that attempts to “tidy up” tags may run counter to “very hook,

attraction, or essence of folksonomies.” The same bind bedevils Wikipedia, the reliability of which has been much debated. Critics of Wikipedia argue that quality is uneven, particularly in specialized topics, and that a lack of respect for expertise in the Wikipedia authoring and editing process is partly to blame for this (see, for example, Lanier, 2006 and associated comments, Sanger, 2004 and associated comments, and Schiff, 2006). While such critics contend that Wikipedia will never become a credible source until its policies change to give greater weight to the opinions of experts, others rebut that such ideas go against the nature of wikis, where the goal is to encourage wide participation (Shirky, 2005a). In Wikipedia’s case, power in its community of encyclopedists is accorded to those with the time and enthusiasm to not only write articles but to respond to the edits of others. If a social classification system were to contemplate some means of explicating semantics, thought would need to be given as to how such explanations would be constructed and who would have the power to do so.

5. Conclusion

Social classification systems have significant advantages, not least of which is the enthusiasm that they engender: users see a benefit to providing metadata, and the process is simple enough that they actually do so. But current implementations of social classification systems also have significant drawbacks that limit their use, some of which are outlined in this paper. One line of research that might proceed from this situation is to improve social classification systems without sacrificing their charm, a challenge outlined by Guy and Tonkin (2006). Another approach is for knowledge organization researchers to take advantage of the current enthusiasm for social classification to better explicate current views of “traditional” classification and its benefits, and to show how other types of knowledge organization schemes may complement use of social classification systems.

Different kinds of knowledge organization schemes are appropriate in different situations. For example, one view is that the diversity of index terms and of term usage in social classification systems promotes the discovery of unanticipated but useful or interesting documents (Kroski, 2005). However, given the unclear semantics for common index terms, the type of discovery enabled by a social classification is more blindly serendipitous than that enabled by a classification scheme with clearly defined concepts and relationships. In contrast, a well-defined classification scheme may be thought of as an argument in service of a particular purpose (according to the philosopher of classification Broadfield, “classification only becomes possible

when I decide what to do with it, and a single definite purpose emerges”) (Broadfield, 1949, p. 87). When this is the case, the classification scheme also promotes discovery of unanticipated documents, but in a way that illuminates a specific interpretation of those documents. One may come upon new and interesting documents via social classification, but it is much more difficult to determine coherent meaning from the structure itself; there is no situating context.

In an article that takes Clay Shirky to task for ignoring the relative utility of different types of knowledge organization schemes, the information architect Peter Merholz pauses in the midst of praising faceted classifications and thesauri to state that, like Shirky, he does “pretty much loathe hierarchy” (Merholz, 2005). As researchers and educators, we should do more to understand the motivation behind such initially puzzling statements and to better articulate the situations in which hierarchy is a preferable strategy, as in Kwasnik, 1999 (and, incidentally, to clarify how faceted classifications and thesauri may also make good use of hierarchy). As a broader goal, along the lines of Jack Andersen’s (2006) examination of information literacy as a complex phenomenon inextricably linked to people’s activities with documents, perhaps we need to consider what it means to be “literate” in the creation of knowledge organization schemes, beyond “standards, techniques, or practical frameworks,” as Andersen says. What do people need to know to be able to effectively use and create systems of organized knowledge, both individually and collaboratively? For example, if, as Andersen claims, information seekers need “genre and activity knowledge” of both documents and of their bibliographic records, do they also need equivalent knowledge regarding the organization *systems* themselves, as structured arguments similar to documents? And if so, what does this mean in terms of creating good information organization systems? Such questions may promote interesting and useful areas of further research.

References

Andersen, J. (2006). The public sphere and discursive activities: Information literacy as sociopolitical skills. *Journal of Documentation* 62(2): 213-238.

Anderson, C. (2004). The long tail. *Wired* 12(10).

Broadfield, A. (1946). *The philosophy of classification*. London: Grafton.

Bruce, H., Jones, W., and Dumais, S. (2004). Information behaviour that keeps found things found. *Information Research* 10(1): paper 207. Retrieved from <http://InformationR.net/ir/10-1/paper207.html>.

Ellis, D. (1989). The behavioral model for information retrieval system design. *Journal of Information Science* 15(4/5): 237-247.

Golder, S., and Huberman, B. (2005). Usage patterns of collaborative tagging systems. *Journal of Information Science* 32(2): 198-208.

Guy, M., and Tonkin, E. (2006). Folksonomies: Tidying up tags? *D-Lib Magazine* 12(1). Retrieved September 14, 2006 from <http://www.dlib.org/dlib/january06/guy/01guy.html>.

Kroski, E. (2005). The hive mind: Folksonomies and user-based tagging. Retrieved September 14, 2006, from <http://infotangle.blogspot.com/2005/12/07/the-hive-mind-folksonomies-and-user-based-tagging/>.

Kwasnik, B. (1999). The role of classification in knowledge representation and discovery. *Library Trends* 48(1): 22-47.

Lanier, J. (2006). Digital Maoism: The hazards of the new online collectivism. *Edge* 183. Retrieved September 14, 2006, from http://www.edge.org/3rd_culture/lanier06/lanier06_index.html.

Lawley, L. (2005a). Social consequences of social tagging. Retrieved September 14, 2006 from http://many.corante.com/archives/2005/01/20/social_consequences_of_social_tagging.php.

Lawley, L. (2005b). It's the social network, stupid. Retrieved September 14, 2006, from http://many.corante.com/archives/2005/01/20/its_the_social_network_stupid.php.

Merholz, P. (2004). Metadata for the masses. Retrieved September 14, 2006, from <http://www.adaptivepath.com/publications/essays/archives/000361.php>.

Merholz, P. (2005). Clay Shirky's viewpoints are overrated. Retrieved September 14, 2006, from <http://www.peterme.com/archives/000558.html>.

Quintarelli, E. (2005). Folksonomies: Power to the people. Paper presented at the ISKO Italy UniMIB meeting, Milan, Italy, June 24, 2005. Retrieved September 14, 2006, from <http://www.iskoi.org/doc/folksonomies.htm>.

Sanger, L. (2004). Why Wikipedia must jettison its anti-elitism. Retrieved September 14, 2006, from <http://www.kuro5hin.org/story/2004/12/30/142458/25>.

Schiff, S. (2006). Know it all: Can Wikipedia conquer expertise? *The New Yorker*, July 31, 2006.

Shirky, C. (2005a). K5 article on Wikipedia anti-elitism. Retrieved September 14, 2006, from http://many.corante.com/archives/2005/01/03/k5_article_on_wikipedia_antielitism.php.

Shirky, C. (2005b). Folksonomies + controlled vocabularies. Retrieved September 14, 2006, from http://many.corante.com/archives/2005/01/07/folksonomies_controlled_vocabularies.php.

Shirky, C. (2005c). Ontologies are overrated: Categories, links, and tags. Retrieved September 14, 2006, from http://www.shirky.com/writings/ontology_overrated.html.

Tuominen, K., Talja, S., and Savolainen, R. (2003). Multiperspective digital libraries: The implications of constructionism for the development of digital libraries. *Journal of the American Society for Information Science and Technology* 54(6): 561–569.

Wilson, P. (1983). *Second-hand knowledge: An inquiry into cognitive authority*. Westport, CT: Greenwood Press.