

Curated Folksonomies: Three Implementations of Structure through Human Judgment

Julia Bullard, School of Information, University of British Columbia

Abstract

Traditional knowledge organization approaches struggle to make large user-generated collections navigable, especially when these collections are quickly growing, in which currency is of particular concern, for which professional classification design is too costly. Many of these collections use folksonomies for labelling and organization as a low-cost but flawed knowledge organization approach. While several computational approaches offer ways to ameliorate the worst flaws of folksonomies, some user-generated collections have implemented a human judgment-centered alternative to produce structured folksonomies. An analysis of three such implementations reveals design differences within the space. This approach, termed “curated folksonomy,” presents a new object of study for knowledge organization and represents one answer to the tension between scalability and the value of human judgment.

1. Introduction

Large, user-generated collections increasingly pose problems for knowledge organization. For platforms such as YouTube and Wikipedia, the collections’ speed of growth outpaces that of traditional organizing schemas, creating collections in which a wealth of content is available and relevant and yet could be functionally invisible (Thornton and McDonald 2012). In designing knowledge organization schemas for large and growing user-generated collections such as these, we accept trade-offs among a method’s organizing functions, its scalability, and its ethical impacts.

These trade-offs become clear as we consider controlled vocabularies, folksonomies, and computational approaches, each of which gives more or less weight to organizing functions, scalability, and ethical impacts. Controlled vocabularies are the gold standard of knowledge organization approaches but require costly expert design and trained indexing labor. They also respond slowly to major changes in the collection, making them ill-suited for rapidly growing, user-generated collections (Yi and Mai Chan 2009; Hoffman 2009; Olson 1998). In contrast, scholars in the field recognize folksonomies as being

deeply-flawed systems that nevertheless make effective use of the distributed, low-effort actions of independent users (Furner 2009; Munk and Mørk 2007; Mai 2011; Golder and Huberman 2006).

As the aggregate of the “personomies” of tag sets users create for largely self-directed purposes (Munk and Mørk 2007), folksonomies represent the diversity of user perspectives (Bates and Rowley 2011), but have many flaws such as ambiguity, variation in granularity, and synonymy (Trant 2009; Golder and Huberman 2006). Attempts to create hybrid systems of “structured folksonomies” (Yoo et al. 2013) that retain the scalability of folksonomies while ameliorating their flaws include linking folksonomy terms to existing controlled vocabularies (Yi and Mai Chan 2009), filtering tags using metrics of consensus (Syn and Spring 2013), and automated mapping into ontologies (Dotsika 2009). These hybrid, computational approaches, although scalable, are not widely implemented (Dotsika 2009) and carry risks such as deferring accountability for harmful outputs to the algorithm or the corpus rather than a designer (Crawford 2016) and further marginalizing minority interpretations by enforcing majority views (Rieder 2016; Aroyo and Welty 2015). Although computational approaches offer means to transform messy folksonomies without expert curation (Zhitomirsky-Geffet et al. 2016), they fail as a solution to dealing with unwieldy collections because they sacrifice human judgment for expediency.

The universe of organizing solutions for user generated collections is not limited to the computational approaches the field has tried so hard to implement. User communities are aware of the flaws of folksonomies and engage in controlled vocabulary design for their own collections. This alternate approach, which I refer to here as a “curated folksonomy,” following the terminology of one particular implementation (Johnson 2014), retains the human judgment of expert controlled vocabulary design while approaching the scalability of computational processes. Curated folksonomies present a possibility for implementing the human judgment of expert controlled vocabulary design at scale by giving users the opportunity to review and revise the folksonomy, interpreting tags as synonymous or related. The same processes that produce large, user-generated collections and their messy folksonomies can also produce controlled vocabularies, given particular system design choices.

The existence of this alternative model indicates a space in knowledge organization for heteromation—a technological approach in which users rather than computers make the critical decisions (Ekbia and Nardi 2014). This model extends and concretizes recurrent ideas in the knowledge organization literature such as democratic indexing (Hilderley and Rafferty 1997), tag gardening (Peters and Weller 2008), and structured folksonomies (Yoo et al. 2013), with a particular emphasis on human decision making in the final form of the knowledge organization system. Exploring the curated folksonomy approach opens up new ways of understanding the concerns and goals of the field of knowledge organization, including notions of power, accountability, and the possibility to represent a plurality of voices.

To introduce this non-computational, hybrid approach, I analyze the design of three sites with current instantiations of curated folksonomies: Stack Overflow, a question and answer site for computer programming; LibraryThing, a social book cataloging site; and Archive of Our Own, a fanwork collection. I examine these curated folksonomies through the lens of traditional knowledge organization criteria of precision, recall, equivalency, hierarchy, and fidelity. My analysis reveals key design choices within the curated folksonomy space that allow for greater complexity of structure or simple solutions to synonymy.

To understand why user communities are adopting curated folksonomies as a knowledge organization approach, and how they differ from computational approaches, it is necessary to explore in more detail how—and how well—folksonomies function. As the raw materials from which curated folksonomies are made, it is important to understand what it is that folksonomies do and how they have illuminated established knowledge organization concerns and concepts. To ground my discussion, I turn to the literature on folksonomies to explain their popularity as knowledge organization systems, their fundamental drawbacks, and the latent conversation in the knowledge organization literature on ways to ameliorate such flaws.

2. Literature Review

Among knowledge organization systems, the folksonomy is notable for its lack of control. In its earliest treatment in this journal, Noruzi (2006) termed the folksonomy an “(un)controlled vocabulary,” and defined it as “an Internet-based information retrieval methodology consisting of collaboratively generated, open-ended labels that categorize content such as web resources, online photographs, and web links” (199). Two elements of this definition are key to distinguishing folksonomies from other modes of knowledge organization: collaborative generation and open-ended labelling. While traditional knowledge organization design often involves elements of collaboration and teamwork, folksonomies begin with the premise that description and retrieval are the cumulative work of a large, distributed set of individuals. Similarly, while some elements of traditional knowledge organization systems account for the inevitability of open-ended fields, this is the premise of a tagging system: that the user will not consult or be limited by a set of predefined terms. These two elements mark out a contrasting space from the dominant paradigm of knowledge organization; Noruzi’s final element, the observation that folksonomies are often used to categorize digital content such as online photographs, begins to explain why these systems have become so ubiquitous in the intervening 12 years. Many contemporary digital collections are user-generated and their size and the speed of their growth defy traditional knowledge organization paradigms.

Folksonomies are ubiquitous across contemporary digital collections but common wisdom and extensive research reveal a few glaring shortcomings (Lee and Schleyer 2012). First, folksonomies suffer from user error, so that misspellings divide otherwise identical tags in retrieval. Second, folksonomies suffer from diversity of user perspectives, so that small differences in grammar (“truck” instead of “trucks”), specificity (“Ford F150” and “truck”), and language (“color” instead of “colour”) produce different retrieval sets. Conversely, folksonomies suffer from convergence through homographs, so that “orange” (color) and “orange” (fruit) are conflated in retrieval. Third, research has found that folksonomies are overwhelmingly populated by self-directed tags, such as “to read,” producing aggregate

sets that serve little function beyond the individual's collection (Golder and Huberman 2006; Munk and Mørk 2007). Despite these contributors to poor precision and recall, folksonomies improve retrieval capability beyond that of full-text searching (Heymann, Koutrika, and Garcia-Molina 2008). They do this by harnessing the energy and interest of users in managing their own objects in instances where more intentional or top-down organization is not feasible.

Design choices in implementing folksonomies can influence these characteristics. For example, the distribution of tags in folksonomies tend to follow “power law” or a long-tail distribution in which some tags are very widely used while the majority of tags have few uses (Munk and Mørk 2007; Golder and Huberman 2006). The tags which do indicate consensus—those on the higher end of the power law distribution—tend to be so superficial and general as to be meaningless in retrieval (Munk and Mørk 2007). Interface design for folksonomies to alleviate expected problems such as synonymy and misspellings, such as producing auto-complete suggestions from the existing folksonomy, risk exacerbating this tendency towards imitation over thoughtfulness (Munk and Mørk 2007). Other design choices, such as computationally deriving suggestions for tags from the items themselves (Razikin et al. 2011) aim to direct taggers to tags more impactful for precision and recall and to reduce the effort of identifying such tags.

Despite the agreed-upon shortcomings of folksonomies, cultural heritage institutions such as libraries, museums, and archives have repeatedly looked to harness their output for retrieval and engagement (Trant 2009; Yi and Mai Chan 2009; Lu, Park, and Hu 2010). Harnessing the output of a folksonomy does not entail ‘fixing’ the folksonomy but instead recognizing what uses and impacts it has beyond those of an expertly-created controlled vocabulary. An exemplar of this approach is Adler's (2014) study of users' tags for transgender-themed books in which Adler noted the importance of open-ended labelling as an empowering discursive practice for individuals and communities historically marginalized and harmed by labelling conventions. Ideally, as user-driven knowledge organization systems, folksonomies provide an opportunity for a community to negotiate terms, express dissent from

the dominant terminology, and otherwise resist and work around more static terminology imposed from the top down.

While we might look to folksonomies as offering something inherently different than traditional knowledge organization systems, the inverse approach is to use existing controlled vocabularies to provide structure and predictability where folksonomies are in need (Yi and Mai Chan 2009; Golub, Lykke, and Tudhope 2014; Dotsika 2009; Matthews et al. 2010). Such approaches recognize the strengths of controlled vocabularies such as their structure and reliability to improve the function of folksonomies to provide access to items. For example, Yi and Chan (2009) mapped the Library of Congress Subject Headings to the Delicious folksonomy. The long-term outcome of this mapping would be to use a controlled vocabulary to enhance information retrieval within a particular folksonomy and across folksonomy collections that have, as a common reference, a relatively stable referent such as LCSH. Similarly, Golub, Lykke, and Tudhope (2014) examined the impact on social tagging decisions when users had suggestions from the Dewey Decimal Classification as they tagged, finding that the influence of the controlled vocabulary was to focus users and lead them to more consistent tags that would better serve the aims of precision and recall.

A third and contrasting approach to either recognizing what folksonomies can add to controlled vocabularies or vice versa is to create hybrid forms in which folksonomies are transformed into structured vocabulary. These approaches go beyond design choices meant to steer users toward more consistent or useful tags (Razikin et al. 2011) but seek to process the folksonomy into a new form. Such approaches might have as their goal a more functional knowledge organization system (Sen et al. 2007; Tsui et al. 2010; Yoo et al. 2013) or, more conservatively, a filter by which expert knowledge organization designers can review folksonomies for relevant suggestions for item description (Syn and Spring 2013). “Democratic indexing” (Brown et al. 1996; Hilderley and Rafferty 1997) was an early version of this approach in which the first pass of user-generated descriptors (roughly equivalent to tags) would undergo a process of “reconciliation” into a public, aggregate view based on computation of the number of times

individuals had used a given descriptor for a given item. Thinking of a folksonomy as raw material from which we might derive a structured knowledge organization system continues to be taken up as a computational problem. In Dotskia's (2009) and Tsui, Wang, Cheung, and Lau's (2010) reviews of this space, folksonomies are seen as amenable to computational approaches such as clustering (Specia & Motta, 2007), matching algorithms (Angeletou et. al, 2007), applying machine-readable dictionaries (Alves, Pereira, & Cardoso, 2002; Rajaraman & Ah-Hwee Tan., 2002), and creating tag networks by a combination of merging and filtering algorithms (Lux & Dosinger, 2007) that can transform a folksonomy into structured systems such as taxonomies and ontologies.

Computational approaches are promising for their scalability; automation reduces the need for ongoing expert labor and many approaches grow more effective with large corpuses and ongoing user interactions (Zhitomirsky-Geffet et al. 2016). However, these approaches carry risks increasingly obvious for computational processing of social tasks and collections: they may defer blame for harmful outputs of the process to the algorithm or the corpus rather than an accountable, human designer (Crawford 2016) and they are likely to further marginalize minorities in favor of enforcing majority views (Aroyo and Welty 2015). Whereas knowledge organization systems are vulnerable to instantiating and reproducing social harms such as discrimination and bias against the less powerful among us (Olson and Schlegl 2002; Berman 1971; Mai 2010), computational approaches may exacerbate these harms, especially where the method is black-boxed or inherently inaccessible to auditing (Ananny and Crawford 2018; Garfinkel et al. 2017; Burrell 2016), and where the machine is entrusted to make the final decision as an objective judge of otherwise messy human processes (Angwin et al. 2016; Noble 2018). For these reasons, it is reasonable for users to be wary of computational approaches to interpreting and reorganizing folksonomies, especially where the community of users is particularly familiar with the flaws of knowledge organization and computational systems.

Here I wish to take up a possibility suggested (Peters and Weller 2008) but not adequately explored in the knowledge organization literature: that the major flaws of a folksonomy might be

addressed using the human judgment of users directly and not through the re-creation of human choices and logic through computational methods. Undoubtedly, folksonomies require some attention given their ubiquity and flaws. Folksonomies are a popular mode of low-cost organization and access for large and otherwise unmanageable user-generation collections that are, in comparison to controlled vocabularies, seriously flawed in their ability to facilitate retrieval across the collection. To be specific, rather than entrusting clustering or algorithmic methods to extract from a folksonomy a sense of term relationships, it is possible to ask users to map these relationships themselves, building off of the work of their tagging to further indicate the structure and equivalencies among the aggregate set of tags. Despite not being a focus of scholarship in this area, this human-focused approach has been adopted and implemented in various forms in online communities. In the next section I will define curated folksonomy as such an approach and detail three such implementations.

3. Curated Folksonomies

I use the term “curated folksonomies” here to collect a few knowledge organization systems currently in use for description and retrieval in large and successful user-generated content sites. Here I will outline the defining characteristics of this form of knowledge organization and explore some of its variant implementations and design choices with particular attention to three instantiations: LibraryThing (librarything.com), Stack Overflow (stackoverflow.com), and Archive of Our Own (archiveofourown.org).

In contrast to classic or unregulated folksonomies, curated folksonomies take the aggregate tags produced by users as a starting point and use expert or collective decision making to identify and alleviate problems with synonymy and homographs. The basic tenets of a curated folksonomy are as follows:

- 1) Users create tags
- 2) Some intentional agent combines synonymous tags and/or differentiates homographic tags

3) Recall and precision are improved

Curated folksonomies are primarily reactive; unlike traditional, top-down controlled vocabulary approaches, most of the terms in a curated folksonomy are driven by user activity, and intentional knowledge organization design follows user action. Curated folksonomies are particularly popular for user-generated collections that are quickly growing, in which currency is of particular concern, for which professional classification design is too costly, and in which users are particularly motivated and suited to engage in organizing work.

Three notable examples of curated folksonomy illustrate these elements. LibraryThing, a user-generated database of books, uses a curated folksonomy to manage tags applied to books, often indicating content, (“animals”), genre (“satire”), or personal relevance (“summer reading”). In LibraryThing’s system of “Tag combining,” all interested users can identify synonymous tags for combining, identify wrongly combined tags for separating, and vote in both types of decisions. The tag combining system is notably strict and follows a relatively narrow sense of synonymy:

Tag combination is driven by a single basic rule: Tags should be combined only when they are the same in both meaning and usage on LibraryThing.

Examples:

- There is no discernible difference in either the use or meaning when it comes to terms like "wwii," "ww2" and "world war two."
- While some might claim they are synonyms, tags like LGBT and GLBT have very different top books. It's likely they encode differences in perspective or identity.¹

In LibraryThing, plurals are not equated to singular versions nor are all abbreviations equated to full terms, in expectation that in some cases these variants actually indicate different meanings. Tags that are

“combined” on LibraryThing redirect to the page for the preferred version of the term. This page displays the collected combined terms and all books tagged with any variant of the term. Usages of variant versions of the tag remain unchanged on book and user pages; combining decisions only affect site-wide retrieval. This design choice keeps intact users’ own organizing schemas within their collections, particularly important for tag variants in other languages (e.g., “ciencia ficcion” and “science fiction”).

The second notable example of an active curated folksonomy is Stack Overflow, a question and answer site for computer programming. All Stack Overflow questions include one or more tags. Tags are used in retrieval of related questions and have information pages (similar to scope notes from traditional controlled vocabulary systems) that give background information. Stack Overflow invites users to “help tame the tag folksonomy.”² Users with higher algorithmic “reputation” on the site can edit tag information pages, can identify tags that are synonymous to “master” tags and vote for or against these connections. That is, Stack Overflow limits the users who can manage the system to those who other users have judged to be good contributors to the community. Tags made synonymous with “master” tags are “automatically and silently remapped” to their master tags; the synonymous forms of tags are changed to the preferred form throughout the site.

These forms of curated folksonomy share the same basic tenets—that users drive the creation of new tags and that synonymous tags are made equivalent for retrieval—while the philosophy of the folksonomies differ. At LibraryThing, tagging is incentivized as a self-directed, personal information management tool: “Once you have a hundred books or so, you need some way to organize them.”³ In this case, the folksonomy is truly “a function of the total sum of persononmies” (Munk and Mørk 2007) or individual tag sets created for the organization and material of a personal collection. Especially with the augmented functionality conferred by the curated folksonomy, the aggregate has the secondary function of supporting retrieval and discovery across the entire site. In contrast, Stack Overflow, as question and answer site, is primarily outward-facing, and tags are intended to help potential answerers monitor

questions in their areas of expertise. The curated folksonomy furthers questioners’ efforts to have their tags seen by relevant answerers by aligning variant tags with expert-preferred synonyms.

The third form of curated folksonomy, from Archive of our Own (AO3), falls in-between the examples of LibraryThing and Stack Overflow. As with LibraryThing, the tagged objects are primarily textual works and the curated folksonomy actions do not change user-chosen variants, only equate terms in retrieval. As with Stack Overflow, the primary activity of the site is outward-facing—authors to readers—and the nominal purpose of tagging and the curated folksonomy is to increase the visibility of user-generated content to relevant readers. Of the three, AO3 is arguably the most selective with regards to which users participate in curated folksonomy design; whereas LibraryThing allows all users to nominate and vote, and Stack Overflow allows established users to nominate and vote, AO3 has a small (~200) team of volunteers who complete a recruitment and training process before receiving access and permissions to the curated folksonomy interface.⁴ Volunteers do not nominate and vote on decisions within the curated folksonomy, rather, each volunteer is responsible for a section of the site and makes the changes to the folksonomy with substantial autonomy.

Table 1 summarizes the variant attributes among classic folksonomies and the systems instantiated within Stack Overflow, LibraryThing, and Archive of Our Own. These attributes represent a starting point for the establishment of curated folksonomies as an object of study within knowledge organization.

	Classic folksonomy	Stack Overflow	LibraryThing	Archive of Our Own
Manual indexing	Yes	Yes	Yes	Yes
Natural language	Yes	Yes	Yes	Yes
Indexing expertise	No	No	No	Training

	Classic folksonomy	Stack Overflow	LibraryThing	Archive of Our Own
Voting	No	Yes	Yes	No
Open participation	Yes	Reputation thresholds	Yes	Application Process
Synonymous term relationships	No	Yes	Yes	Yes
Differentiate homographs	No	No	No	Yes
Hierarchical term relationships	No	No	No	Yes
Preserves user variants	Yes	No	Yes	Yes

Table 1: Attributes of Curated Folksonomy Instantiations

In combination, these attributes address the major flaws of folksonomies that impede precision and retrieval. Only one of these versions provides a mechanism to discriminate between homographs while all of the curated folksonomy versions take as their focus the ability to equate synonyms. Indicating a more substantial alternative to the complex computational approaches that create taxonomic or ontological structures (as in Zeng 2008), the instantiation at Archive of Our Own also creates hierarchical relationships among terms. Noting the existence of curated folksonomies as a variant knowledge organization approach is a preliminary step. In the following section I suggest productive avenues of inquiry for knowledge organization research.

4. Next Steps

Two broad categories of response follow from the existence and the arguable success of curated folksonomies in addressing the major shortcomings of folksonomies and providing organization and retrieval within large, user-generated collections. The first is how we might study the curated folksonomy as a variant form of knowledge organization. The second is what the existence of curated folksonomies

and their place among knowledge organization forms suggests are meaningful questions within knowledge organization more broadly.

4.1 Studying Curated Folksonomies

The first set of questions we might ask of curated folksonomies and their variant forms is whether they work, within the set of evaluative criteria established as meaningful within knowledge organization. That is, do they facilitate precision and recall in retrieval (Spärck Jones 2005)? This might be asked and answered in at least two ways: Do they achieve a measurable improvement over classic folksonomies in this regard and do they approach anything like the standard established by controlled vocabularies? Do they create a set of terms and term-object relationships consistent with the values of the community (Bullard 2017; Feinberg 2007; Mai 2010)? Is the resulting system ethically sound and defensible within a recognized ethical paradigm (Fox and Reece 2012)? Are the equivalency and hierarchical relationships upon which users arrive logically sound (Frické 2016; Furner 2012)? The answers to these questions will differ across variants of the form and according to the needs and values of the communities and collections to which they are applied. Given that the strength of the curated folksonomy approach is its applicability to large and growing user-generated collections, one particularly relevant evaluative criterion is the sustainability of these systems over time and at scale (Ibekwe-Sanjuan and Bowker 2017). Notably, each of three examples I have chosen to illustrate the curated folksonomy approach come from communities that self-select for individuals interested in manipulating specialized vocabularies, organizing items, or creating new relationships among works, leaving open the possibility that their success might not be as easily reproduced in community collections without such relevant characteristics.

Alongside these questions of how curated folksonomies might fare in our evaluative constructs is how this form of knowledge organization takes up, modifies, and possibly complicates established theories and processes from knowledge organization theory. As a primarily reactive design method, contrasting most sharply with traditional top-down and expert-driven knowledge organization design, curated folksonomies may invite a different mode of evaluation. Rather than focusing on the product, that

is, the resulting system, its characteristics and function, curated folksonomies may be particularly appropriate for evaluative methods focusing on process. That is, are the methods of voting, negotiation, or authority in transforming the folksonomy into a set of term relationships consistent with prescriptive models of knowledge organization, such as the correct application of hierarchical relationships? Would some variants of the curated folksonomy form actually be more appropriately “democratic” in politics as opposed to the libertarian nature of classic folksonomies (Feinberg 2006)?

Finally, we might turn the analysis from the functions within these systems to the function they play in context of communities and collections. In summarizing the current state of folksonomies and scholarly responses to their function and prevalence, I argued that discomfort with computational approaches over issues of accountability and transparency are sufficient for users to be wary over their application to the construction of controlled vocabularies. I base this claim on contemporary critiques of algorithmic systems and the established vulnerability of knowledge organization systems to perpetuating discrimination and bias. This explanation for the adoption of more human-focused methods such as curated folksonomies requires further interrogation and testing. The three communities with notable curated folksonomy instantiations are particularly aware of these vulnerabilities: LibraryThing as a contrast to traditional library systems and the colonial and imperialistic histories of the Library of Congress and Dewey systems, Stack Overflow as a community based on the fact that programming is inevitably full of errors, and Archive of Our Own as an explicitly intersectional feminist community accustomed to being mislabeled by dominant cultures (Fiesler, Morrison, and Bruckman 2016). It is worth asking whether these qualities, particular to these three communities, are related to the adoption of a human-focused system and whether other motivations and causes are at play in the choice to take on additional human labor as an alternative to keeping a classic folksonomy or implementing a computational remedy.

4.2 Implications for Knowledge Organization

As with any expansion of the scholarly space, to fit curated folksonomies into knowledge organization not only provides a set of questions to ask of this new form but suggests activities to readjust the space itself in response. Particularly because this form of knowledge organization originates from outside the established professional community, this opportunity comes years after these systems have been in implementation. Here I suggest a few questions we might have opportunity to revisit in the process of integrating curated folksonomies among the various characteristics of knowledge organization systems. First, as an intersection of bottom-up and top-down design methods, curated folksonomies suggest new ways to understand the convergence and divergence between the philosophies of these two extremes (Mai 2011; Feinberg 2006; Furner 2009). Second, as an alternative to computational approaches, curated folksonomies embody a validation of human judgment and domain expertise, suggesting a possible alignment with knowledge organization approaches that value the human as an instrument of analysis, interpretation, and accountability (Hjørland and Albrechtsen 1995; Feinberg 2011). Similarly, curated folksonomies might be counted among other indicators that knowledge organization—particularly forms centering human as opposed to machine judgment—is on its pendulum swing away from universal, totalizing systems and back toward local and specialized systems (Smiraglia 2002; Augusto et al. 2016). Finally, the reactive design of curated folksonomies invites a shift in focus from the initial design and implementation of knowledge organization systems to maintenance and revision as central practices. While the adaptation of existing knowledge organization systems is a matter of perennial interest (Olson 1998; Nero 2006), and hospitality to expansion is a fundamental characteristic of many knowledge organization forms, maintenance and revision are seldom presented as typical or central modes of knowledge organization work (Soergel 1974; Park 2008).

Additionally, considering new forms of knowledge organization such as curated folksonomies invites us to look outward at cognate fields and applications. Given that curated folksonomies are inherently collaborative forms with particular technological needs around participation and responsiveness between the folksonomy and its revisions, the study of this form requires connections with

the scholarship of computer-supported cooperative work, crowdsourcing, and participatory design. As a contrast to computational approaches, this form also suggests a growing space for knowledge organization scholars within the contemporary discussion of heteromation (Ekbja and Nardi 2014) and resistance to automation and algorithmic governance (Zarsky 2016; Noble 2018).

5. Conclusion

Large online collections pose challenges for knowledge organization. Ceding the ground to folksonomies is an easy out that leaves these collections and their users with the bare minimum of precision and recall capabilities. Improving folksonomies through computational approaches that harness the size and growth of the vocabulary set is an appealing approach under investigation within knowledge organization and computer science scholarship. Here I presented an alternative approach implemented by three popular user-generated collections that instantiates theories and recommendations latent in the knowledge organization literature. In particular, the curated folksonomy approach emphasizes human judgment in the contemporary context in which vulnerable communities are wary of computational decision making.

Curated folksonomies appear to have grown ‘wild’ outside of the traditional, professional domain of knowledge organization. Bringing curated folksonomies within this scholarly domain can improve the systems themselves; the design variations among the three implementations here indicate a possible range of choices that can impact the systems’ accuracy, functionality, and scalability. As with any knowledge organization approach, curated folksonomies are subject to analysis and evaluation with regards to established attributes such as precision and recall, logic, and representation. Bringing curated folksonomies within our domain also means making space for solutions developed in the wild in possible opposition to established and dominant trends within knowledge organization: distributed rather than centralized design, user rather than professional control, and human rather than computational processing.

Notes

¹ https://wiki.librarything.com/index.php/Tag_combining. Wiki page last edited August 4, 2018, accessed November 9, 2018.

² <https://stackoverflow.blog/2010/08/01/tag-folksonomy-and-tag-synonyms/>. Blog post August 1, 2010, accessed November 9, 2018.

³ <http://www.librarything.com/concepts#what>. Undated webpage, accessed November 9, 2018.

⁴ https://archiveofourown.org/faq/tags?language_id=en. Undated webpage, accessed November 9, 2018.

References

- Ananny, Mike, and Kate Crawford. 2018. "Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability." *New Media & Society* 20 (3): 973–89. <https://doi.org/10.1177/1461444816676645>.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. "Machine Bias." *ProPublica*, May 23, 2016.
- Aroyo, Lora, and Chris Welty. 2015. "Truth Is a Lie: Crowd Truth and the Seven Myths of Human Annotation." *AI Magazine* 36 (1): 15–24. <https://doi.org/10.1609/aimag.v36i1.2564>.
- Augusto, José, Chaves Guimarães, Fabio Assis Pinho, and Suellen Oliveira Milani. 2016. "Theoretical Dialogs About Ethical Issues in Knowledge Organization: García Gutiérrez, Hudon, Beghtol, and Olson." *Knowledge Organization* 43 (5): 338–51.
- Bates, Jo, and Jennifer Rowley. 2011. "Social Reproduction and Exclusion in Subject Indexing." *Journal of Documentation* 67 (3): 431–48. <https://doi.org/10.1108/00220411111124532>.
- Berman, Sanford. 1971. *Prejudices and Antipathies: A Tract on the LC Subject Heads Concerning People*. Metuchen, New Jersey: Scarecrow Press.
- Brown, Pauline, Rob Hilderley, Hugh Griffin, and Sarah Rollason. 1996. "The Democratic Indexing of Images." *New Review of Hypermedia and Multimedia* 2 (1): 107–20. <https://doi.org/10.1080/13614569608914677>.
- Bullard, Julia. 2017. "Warrant as a Means to Study Classification System Design." *Journal of Documentation* 73 (1): 75–90. <https://doi.org/10.1108/JD-06-2016-0074>.
- Burrell, Jenna. 2016. "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms." *Big Data & Society* 3 (1): 205395171562251.

<https://doi.org/10.1177/2053951715622512>.

Crawford, Kate. 2016. "Can an Algorithm Be Agonistic? Scenes of Contest in Calculated Publics." *Science, Technology & Human Values* 41 (1): 77–92.

Dotsika, Fefie. 2009. "Uniting Formal and Informal Descriptive Power: Reconciling Ontologies with Folksonomies." *International Journal of Information Management* 29 (5): 407–15.
<https://doi.org/10.1016/j.ijinfomgt.2009.02.002>.

Ekbia, Hamid, and Bonnie Nardi. 2014. "Heteromation and Its (Dis)Contents: The Invisible Division of Labor between Humans and Machines." *First Monday* 19 (6): 1–15.
<https://doi.org/10.5210/fm.v19i6.5331>.

Feinberg, Melanie. 2006. "An Examination of Authority in Social Classification Systems." *Advances in Classification Research Online*, no. 1: 1–11.
<https://doi.org/10.7152/acro.v17i1.12490>.

———. 2007. "Hidden Bias to Responsible Bias: An Approach to Information Systems Based on Haraway's Situated Knowledges." *Information Research* 12 (4).

———. 2011. "How Information Systems Communicate as Documents: The Concept of Authorial Voice." *Journal of Documentation* 67 (6): 1015–37.
<https://doi.org/10.1108/00220411111183573>.

Fiesler, Casey, Shannon Morrison, and Amy S Bruckman. 2016. "An Archive of Their Own: A Case Study of Feminist HCI and Values in Design." *CHI '16 Proceedings of the ACM Conference on Human Factors in Computing Systems*.
<https://doi.org/10.1145/2858036.2858409>.

Fox, Melodie J., and Austin Reece. 2012. "Which Ethics? Whose Morality?: An Analysis of Ethical Standards for Information Organization." *Knowledge Organization* 39 (5): 377–83.

- Frické, Martin. 2016. "Logical Division." *Encyclopedia of Knowledge Organization*.
International Society for Knowledge Organization.
- Furner, Jonathan. 2009. "Folksonomies." *Encyclopedia of Library and Information Sciences, Third Edition*, 1858–66. Editors Marcia J. Bates and Mary Niles Maack. Boca Raton: CRC Press.
- . 2012. "FRSAD and the Ontology of Subjects of Works." *Cataloging & Classification Quarterly* 50 (5–7): 494–516. <https://doi.org/10.1080/01639374.2012.681269>.
- Garfinkel, Simson, Jeanna Matthews, Stuart S. Shapiro, and Jonathan M. Smith. 2017. "Toward Algorithmic Transparency and Accountability." *Communications of the ACM* 60 (9): 5–5. <https://doi.org/10.1145/3125780>.
- Golder, Sa, and Ba Huberman. 2006. "Usage Patterns of Collaborative Tagging Systems." *Journal of Information Science* 32 (August 2005): 198–208. <https://doi.org/10.1177/0165551506062337>.
- Golub, Koraljka, Marianne Lykke, and Douglas Tudhope. 2014. "Enhancing Social Tagging with Automated Keywords from the Dewey Decimal Classification." *Journal of Documentation* 70 (5): 801–28. <https://doi.org/10.1108/JD-05-2013-0056>.
- Heymann, Paul, Georgia Koutrika, and Hector Garcia-Molina. 2008. "Can Social Bookmarking Improve Web Search?" In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 195–206. WSDM '08. New York, NY, USA: ACM. <https://doi.org/10.1145/1341531.1341558>.
- Hidderley, Rob, and Pauline Rafferty. 1997. "Democratic Indexing: An Approach to the Retrieval of Fiction." *Information Services & Use* 17 (2–3): 101–9. <https://doi.org/10.3233/ISU-1997-172-304>.

- Hjørland, Birger, and Hanne Albrechtsen. 1995. "Toward a New Horizon in Information Science: Domain-Analysis." *Journal of the American Society for Information Science* 46 (6): 400–425.
- Hoffman, GL. 2009. "Applying the User-Centered Paradigm to Cataloging Standards in Theory and Practice: Problems and Prospects." *NASKO* 2: 27–34.
- Ibekwe-Sanjuan, Fidelia, and Geoffrey C. Bowker. 2017. "Implications of Big Data for Knowledge Organization." *Knowledge Organization* 44 (3): 187–99.
- Johnson, Shannon Fay. 2014. "Fan Fiction Metadata Creation and Utilization within Fan Fiction Archives: Three Primary Models." *Transformative Works and Cultures* 17 (2014): 1–15. <https://doi.org/10.3983/twc.v17i0.578>.
- Lee, Danielle H., and Titus Schleyer. 2012. "Social Tagging Is No Substitute for Controlled Indexing: A Comparison of Medical Subject Headings and CiteULike Tags Assigned to 231,388 Papers." *Journal of the American Society for Information Science and Technology* 63 (9): 1747–57. <https://doi.org/10.1002/asi.22653>.
- Lu, Caimei, J.-r. Park, and Xiaohua Hu. 2010. "User Tags versus Expert-Assigned Subject Terms: A Comparison of LibraryThing Tags and Library of Congress Subject Headings." *Journal of Information Science* 36 (6): 763–79. <https://doi.org/10.1177/0165551510386173>.
- Mai, Jens-Erik. 2010. "Classification in a Social World: Bias and Trust." *Journal of Documentation* 66 (5): 627–42. <https://doi.org/10.1108/00220411011066763>.
- . 2011. "Folksonomies and the New Order: Authority in the Digital Disorder." *Knowledge Organization* 38 (2): 114–22.
- Matthews, Brian, Catherine Jones, Bartłomiej Puzon, Jim Moon, Douglas Tudhope, Koraljka

- Golub, and Marianne Lykke Nielsen. 2010. "An Evaluation of Enhancing Social Tagging with a Knowledge Organization System." *ASLIB Proceedings* 62 (4/5): 447–65.
<https://doi.org/10.1108/00012531011074690>.
- Munk, T Bisgaard, and Kristian Mørk. 2007. "Folksonomy, the Power Law & the Significance of the Least Effort." *Knowledge Organization* 34 (1): 16–33.
- Nero, LM. 2006. "Classifying the Popular Music of Trinidad and Tobago." *Cataloging & Classification Quarterly* 42 (3–4): 119–33. <https://doi.org/10.1300/J104v42n03>.
- Noble, Safiya. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.
- Noruzi, Alireza. 2006. "Folksonomies: (Un)Controlled Vocabulary?" *Knowledge Organization* 33 (4): 199–203.
- Olson, Hope A. 1998. "Mapping beyond Dewey's Boundaries: Constructing Classificatory Space for Marginalized Knowledge Domains." *Library Trends* 47 (2): 3–20.
- Olson, Hope A, and Rose Schlegl. 2002. "Standardization, Objectivity, and User Focus: A Meta-Analysis of Subject Access Critiques." *Cataloging & Classification Quarterly* 32 (2): 61–80. https://doi.org/10.1300/J104v32n02_06.
- Park, Ok Nam. 2008. "Opening Ontology Design: A Study of the Implications of Knowledge Organization for Ontology Design." *Knowledge Organization* 35 (4): 209–21.
- Peters, Isabella, and Katrin Weller. 2008. "Tag Gardening for Folksonomy Enrichment and Maintenance." *Webology* 5 (3).
- Razikin, K., D. H. Goh, A. Y. K. Chua, and Chei Sian Lee. 2011. "Social Tags for Resource Discovery: A Comparison between Machine Learning and User-Centric Approaches." *Journal of Information Science* 37 (4): 391–404.

<https://doi.org/10.1177/0165551511408847>.

Rieder, Bernhard. 2016. "Scrutinizing an Algorithmic Technique: The Bayes Classifier as Interested Reading of Reality." *Information, Communication & Society*, April, 1–18.

<https://doi.org/10.1080/1369118X.2016.1181195>.

Sen, Shilad, F. Maxwell Harper, Adam LaPitz, and John Riedl. 2007. "The Quest for Quality Tags." In *Proceedings of the 2007 International ACM Conference on Conference on Supporting Group Work - GROUP '07*, 361. New York, New York, USA: ACM Press.

<https://doi.org/10.1145/1316624.1316678>.

Smiraglia, Richard P. 2002. "The Progress of Theory in Knowledge Organization." *Library Trends* 50 (3): 330–49.

Soergel, Dagobert. 1974. *Indexing Languages and Thesauri: Construction and Maintenance*. Los Angeles: Melville Publishing Company.

Spärck Jones, Karen. 2005. "Some Thoughts on Classification for Retrieval." *Journal of Documentation* 61 (5): 571–81. <https://doi.org/10.1108/00220410510625796>.

Syn, SY, and MB Spring. 2013. "Finding Subject Terms for Classificatory Metadata from User-Generated Social Tags." *Journal of the American Society for Information Science and Technology* 64 (5): 964–80. <https://doi.org/10.1002/asi>.

Thornton, Katherine, and David W. McDonald. 2012. "Tagging Wikipedia: Collaboratively Creating a Category System." In *Proceedings of the 2012 ACM Conference on Supporting Group Work*, 219–27. <http://dl.acm.org/citation.cfm?id=2389210>.

Trant, Jennifer. 2009. "Studying Social Tagging and Folksonomy: A Review and Framework." *Journal of Digital Information* 10 (1).

Tsui, Eric, W. M. Wang, C. F. Cheung, and Adela S M Lau. 2010. "A Concept-Relationship

Acquisition and Inference Approach for Hierarchical Taxonomy Construction from Tags.”

Information Processing and Management 46 (1): 44–57.

<https://doi.org/10.1016/j.ipm.2009.05.009>.

Yi, Kwan, and Lois Mai Chan. 2009. “Linking Folksonomy to Library of Congress Subject

Headings: An Exploratory Study.” *Journal of Documentation* 65 (6): 872–900.

<https://doi.org/10.1108/00220410910998906>.

Yoo, D., K. Choi, Y. Suh, and G. Kim. 2013. “Building and Evaluating a Collaboratively Built

Structured Folksonomy.” *Journal of Information Science* 39 (5): 593–607.

<https://doi.org/10.1177/0165551513480309>.

Zarsky, T. 2016. “The Trouble with Algorithmic Decisions: An Analytic Road Map to Examine

Efficiency and Fairness in Automated and Opaque Decision Making.” *Science, Technology*

& Human Values 41 (1): 118–32. <https://doi.org/10.1177/0162243915605575>.

Zeng, Marcia Lei. 2008. “Knowledge Organization Systems (KOS).” *Knowledge Organization*

35 (2/3): 160–82.

Zhitomirsky-Geffet, M., B.H. Kwaśnik, J. Bullard, L. Hajibayova, J. Hamari, and T. Bowman.

2016. “Crowdsourcing Approaches for Knowledge Organization Systems: Crowd

Collaboration or Crowd Work?” *Proceedings of the Association for Information Science*

and Technology 53 (1). <https://doi.org/10.1002/pa2.2016.14505301013>.