

Complex Decisions: Overturning the Rules to Create Inclusive Systems

Julia Bullard, University of British Columbia

In this piece I discuss how workers in an atypical knowledge organization system managed conflicts between rules and what was best for users. Primarily, this conflict occurred when the workers identified ways in which the system's priorities of precision, recall, and ease of use would disproportionately harm a minority of users. The decisions of these workers were grounded in the context in which they made these decisions—the controlled vocabulary for a user-generated fanfiction collection—but touch on concerns common to workers across a variety of forms of knowledge organization systems.

The workers I discuss here differed from the typical knowledge organization worker not only by virtue of the content of their collection but also in that they were volunteers and required no formal credentials, though they received a few months' worth of part-time training. They shared in common with cataloguers and indexers in more traditional contexts a primary goal to enable user access, a concern with preferred terminology, a focus on text retrieval, and that many of them were credentialed librarians and library technicians. In comparing the perspective of these workers to other knowledge organization contexts I found familiar concerns with layers of technical infrastructure and awareness of the impact of terminology choices on the users' experience of fair and just service.

My aim in studying the daily tasks of these workers was to determine the types of information and outcomes one considers when making decisions of terms and term relationships for organizing items in a collection. There has long been an understanding that our systems enact and perpetrate harmful bias (Bates & Rowley, 2011; Berman, 1971; Olson, 2001). Two approaches have resulted from this understanding: first, to examine our existing systems to identify the types of harmful bias they enact and suggest revisions to remedy them (e.g., Olson, 1998) and second, to accept the inevitability that our systems will be biased and to establish trust with users through transparency and meaningful choices around what type of bias we enact (Feinberg, 2007; Mai, 2010). A complementary approach, which I pursue here, is to look closely at system design and the choices available to workers as they construct, revise, and maintain knowledge organization systems with an aim to inform responsible design practices and training.

To study how knowledge organization systems come to take on their particular flavour of bias and what role worker decisions have in the final character of the system, I approached a large and growing system through a participant observation method; I volunteered to assist with system maintenance in exchange for access to privileged system information and to recruit from among the workers for diary studies and interviews. Nineteen workers, of the 200 active volunteers, completed two-week diaries detailing their daily tasks, followed by 90-120 minute interviews expanding on their reports. Of particular focus in the diary and interview questions were "complex tasks," which the workers were free to define for themselves. In brief, what differentiated simple from complex tasks in their daily work was how many sources of information they would have to consult to be confident of their decision and how confident they could ultimately be that their decision was the right one.

I identified ten factors common to the workers' reflections on complex decisions. These factors represent the types of concerns workers had in determining correct terms and the relationships among terms; the most complex decisions were those in which these factors conflicted, giving the workers alternate and contradictory direction. Four of these factors are foundational to an organizing system of terms and should be familiar despite the particular form (such as classification scheme, catalogue, index, thesauri) that the system takes: ambiguity, filtering, hierarchy, and temporality.



In other words, workers first considered whether the choice of a term was sufficiently differentiated from other meanings, that the term would allow users to narrow the collection to appropriate items, that the broader- and narrower-term relationships were accurate, and that the system would continue to be accurate and usable over time. Two further factors, user primacy and the user/designer gap, are similarly common to user-centered design: the desire to produce a system primarily accountable to the users' needs and the understanding that the user's perspective of the system is inevitably different than that of the designer.

A third set of factors, autocomplete and the server indexing burden, were specific to the technical implementation of this particular system. These factors were relevant to the workers as they considered how their actions would interact with the interface and operation of the site as a whole. For example, workers knew that if they designated a term as a preferred term in the system, it would appear as an autocomplete suggestion for users as they entered an overlapping string of characters into a search field. The worker might also reconsider making a change to a term when the resulting changes to indexed items throughout the system would tie up the server for hours, making the collection slow to respond to user requests. Though these two factors were specific to this technical implementation, a worker in another context might imagine similar limitations and interactions that shape their decisions for terms and records: the maximum character length of a field, the required labour to relabel or reshelve corrected items, or the system's ability to connect letter variants to simplified forms (e.g., ç and c).

As an example of how these factors impacted workers' decisions on terms, consider how users might interpret the suggestion from autocomplete, as they enter the name "Dumbledore," that they complete the descriptor as "Dumbledore Lives." Fanfiction works typically take as their starting point a divergence from source material plot points; therefore, for an experienced fanfiction reader, the presence of "Dumbledore Lives" as a descriptor for fanfiction implies that Dumbledore, in the source material, dies. Particularly when users' familiarity with the source material varies according to their engagement with book or theatrical release schedules, interacting with the system can spoil plot points of which they were previously (and intentionally) unaware. In making a decision to designate a term such as "Dumbledore Lives" as a preferred term, after taking into account other factors, including the prominence the descriptor across item records, the worker might overturn an otherwise sound decision in order to spare a few users an ill-timed spoiler. In the words of one worker from this study, *"I do spend an inordinate amount of time trying not to make teenage girls cry."* The point here is to highlight that, though these tasks are ultimately about connecting users to items through the construction of a record, our decisions impact the system in other ways.

Complex decisions occurred when the rules did not account for the workers' on-the-ground understanding of the collection and user needs. As in the previous example, breaking the rules or sacrificing the system's precision, recall, and ease of use often meant improving the experience of the service or collection for a small minority of users (e.g., those participating in a fanfiction community around source material they have not fully consumed). The rules enacted a utilitarian ethic: they designated choices that would do the most good (largely in terms of improving precision, recall, and ease of use) for the most users (Fox & Reece, 2012; Mills, 1979). In this system, the workers were given a great deal of autonomy in adapting the universal rules for their specific areas of responsibility and a common reason to diverge from those rules was to improve the experience of a small minority of users. In some instances, as when deciding how some subjects were grouped as "folklore" or "mythology" as opposed to "religion," the minority the workers were concerned with were those who had experienced historical oppression. In other instances, the minority was defined within the fanfiction community, such as when the workers were concerned for the function of the collection among users who followed a less popular character or story that was typically dominated in traffic, users, and items by a more widely-distributed and consumed version. These types of decisions, respectively, accounted for two factors common to complex decisions in the workers' reports: authenticity, or the importance of representing concepts such as identity consistently with members of those communities, and inclusivity, or the importance of creating a terminology system that was welcome to diverse users and that did not enact ongoing forms of harm and oppression.

Some decisions we might think of as being complex, the workers in this study did not identify as such. The issue of how to represent changes in an individual's name, particularly around gender identity, is not yet fully resolved in cataloguing (Thompson, 2016). This is not only a matter of whether the individual's chosen name is a preferred term or the main entry in an authority record. Catalogue records may also include inaccurate pronouns and make prominent an individual's prior name (a potentially unwelcome association). Among the fanfiction community—which typically shares an inclusive and progressive ethos (Fiesler, Morrison, & Bruckman, 2016; Hellekson & Busse, 2006; Katyal, 2006)—the ability of an individual to determine their own name, particularly when a name is an expression of gender, was so commonly accepted as a basic human right that all other relevant factors (such as user familiarity or legal 'correctness') was suspended, without trouble, in order to designate a name accordingly. This decision was not 'complex' because the workers had already decided the issue conclusively; regardless of all other factors, a transgender individual would always be known by the individual's chosen name. In fact, the only incident within my study in which an individual's choice of their own name caused any conflict was when a worker had incomplete information (because she had not been keeping up with the relevant television show) to know that a character had chosen a new name to match their gender identity. It was not until a user complaint pointed out the discrepancy that the worker changed the designated name to that chosen by the character; there was no reluctance on the part of the worker to revise the record, only a limitation of the currency of their domain knowledge.

My study of the daily tasks of these workers highlighted the importance of thoughtful, compassionate human judgment in knowledge organization. Given a wealth of rules and procedures, a shifting set of technical requirements and layers of infrastructure, and a large and diverse set of users, knowledge organization workers are in the position to define the character and bias of a system through careful interpretation and exemptions to the rules. With more transparent accounts of the daily processes of system maintenance and revision we might expand and test these factors.

Works Cited

- Bates, J., & Rowley, J. (2011). Social reproduction and exclusion in subject indexing. *Journal of Documentation*, 67(3), 431–448. <http://doi.org/10.1108/00220411111124532>
- Berman, S. (1971). *Prejudices and antipathies: A tract on the LC subject heads concerning people*. Metuchen, New Jersey: Scarecrow Press.
- Feinberg, M. (2007). Hidden bias to responsible bias: an approach to information systems based on Haraway's situated knowledges. *Information Research*, 12(4).
- Fiesler, C., Morrison, S., & Bruckman, A. S. (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*. <http://doi.org/10.1145/2858036.2858409>
- Fox, M. J., & Reece, A. (2012). Which ethics? Whose morality?: An analysis of ethical standards for information organization. *Knowledge Organization*, 39(5), 377–383.

- Hellekson, K., & Busse, K. (2006). *Fan fiction and fan communities in the age of the internet*. McFarland.
- Katyal, S. K. (2006). Performance, property, and the slashing of gender in fan fiction. *The American University Journal of Gender, Social Policy & the Law*, 14(3), 461–518. <http://doi.org/10.1525/sp.2007.54.1.23>.
- Mai, J.-E. (2010). Classification in a social world: bias and trust. *Journal of Documentation*, 66(5), 627–642. <http://doi.org/10.1108/00220411011066763>
- Mills, J. S. (1979). *Utilitarianism*. (G. Sher, Ed.). Indianapolis: Hackett Pub. Co.
- Olson, H. A. (1998). Mapping beyond Dewey's boundaries: Constructing classificatory space for marginalized knowledge domains. *Library Trends*, 47(2), 3–20.
- Olson, H. A. (2001). The Power to Name: Representation in Library Catalogs. *Signs*, 26(3), 639–668.
- Thompson, K. (2016). More than a name: A content analysis of name authority records for authors who self-identify as trans. *Library Resources & Technical Services*, 6(3). <http://doi.org/10.5860/lrts.60n3.140>