The current issue and full text archive of this journal is available at
www.emeraldinsight.com/0022-0418.htm

Classification in a social world: bias and trust

Jens-Erik Mai
Faculty of Information, University of Toronto, Toronto, Canada

Received 5 November 2009
Revised 14 December 2009
Accepted 15 December 2009

Abstract

Purpose – The purpose of this paper is to establish pluralism as the basis for bibliographic classification theory and practice and examine the possibility of establishing trustworthy classifications.

Design/methodology/approach – The paper examines several key notions in classification and extends previous frameworks by combining an explanation-based approach to classification with the concepts of cognitive authority and trust.

Findings – The paper presents an understanding of classification that allows designers and editors to establish trust through the principle of transparency. It demonstrates that modern classification theory and practice are tied to users' activities and domains of knowledge and that trustworthy classification systems are in close dialogue with users to handle bias responsible and establish trust.

Originality/value – The paper establishes a foundation for exploring trust and authority for classification systems.

Keywords Classification, Philosophical concepts, Epistemology, Trust, Authority

Paper type Conceptual paper

Introduction

Jorge Luis Borges wonders in his essay "The analytic language of John Wilkins" why we have attempted to describe the world, as the world really is, when "obviously there is no classification of the universe that is not arbitrary and conjectural" (Borges, 1952, p. 104). While most people might nod in agreement with Borges' statement, most bibliographic classifications appear to take a different stand. The world of information has changed radically recently; it seems that the world of information is more social, more about interactions, collaboration, sharing, and more fluid. Whereas classification systems traditionally have exhibited authority by virtue of their status and control, many of them are being questioned today and bibliographic classification need to find ways to deal with charges of bias and find new ways to establish trust with their users (Beghtol, 1986).

Borges (1952, p. 103) gives as an example of a classification:

> [. . .] a certain Chinese encyclopedia entitled "Celestial Empire of benevolent Knowledge". In its remote pages it is written that the animals are divided into: (a) those that belong to the emperor, (b) embalmed ones, (c) those that are trained, (d) sucking pigs, (e) mermaids, (f) fabulous ones, (g) stray dogs, (h) those that are included in this classification, (i) those that tremble as if they were mad, (j) innumerable ones, (k) those drawn with a very fine camel's hair brush, (l) others, (m) those that have just broken a flower vase, (n) those that resemble flies from a distance.

Bowker and Star (1999, p. 131) use Borges' Chinese animal classification as an example to demonstrate that many, perhaps all, classifications at the surface may look innocent and perfectly fine, but that on close inspections all classification reveal real

Journal of Documentation
Vol. 66 No. 5, 2010
pp. 627-642

0022-0418
DOI 10.1108/00220411011066763

consequences and assumptions about the world; however, classifications might "as with many strange things, [. . .] become well adopted to the modern bureaucracy". To understand the basic assumptions hidden in classifications, Bowker and Star (1999, p. 131) argue, we need to bring classifications out of their contexts, "classifications that appear natural, eloquent, and homogenous within a given human context appear forced and heterogeneous outside that context"; only outside their contexts will the forced nature of classifications be revealed. The arbitrary and conjectural nature of all classifications is revealed when they are taken outside their context and as such there the contextuality of classifications are emphasized.

Shirky (2005), in his popular essay "Ontology is overrated: categories, links, and tags," argues that classifications like the A section (Marxism-Leninism) in the former Soviet classification system, Bibliotečno-Bibliografičeskaja Klassifikacija, the 200s (Religion) in Dewey decimal classification (DDC), and section D (History (general)) in the Library of Congress Classification (LCC) are biased, whereas the periodic table of the elements, according to Shirky (2005), is the "Best. Classification. Evar" because it has "both descriptive and predictive value"[1]. The basic assumption that Shirky makes is that some classifications are not as good as others because they are biased; in opposition to that view, Feinberg (2007) argues that, in fact, all classifications are biased and that bias is not a feature of a classification that can be eliminated and erased to create a neutral and unbiased classification. She further argues that, "if we cannot eliminate bias, then we should instead attempt to be more responsible about bias and explicitly decide on and defend the perspectives represented in our information systems" (Feinberg, 2007). It seems that the qualities and objectives of classification need to be reconsidered, revised, and rediscovered – concepts such as neutrality, bias, and authority are under pressure and force a reconceptualization of classification.

Classification in library and information science has long been occupied with finding the one perfect system; the system that works for everyone, everywhere (Miksa, 1998). However, basic work and thinking about categories and classification has long told us that achieving that one perfect system is not possible (Lakoff, 1987, p. 12ff) and recent developments in social technologies have challenged that goal even further (Weinberger, 2007; Shirky, 2008). In this paper, I will explore the conceptual foundations of classification in a social world and present an understanding for working with classification in library and information science that rest upon a social understanding of how knowledge and information is achieved, generated, and trusted. Following Bryant (2000), Wilson (1983) and Broadfield (1946), this explanation-based approach to classification rests on the main notion that reality does not exist in a knowable form and that classifications need to be couched in an explanation for users to trust them.

The one best system

Harris (1870) described the challenges of librarians – and anyone else charged with organizing and retrieving recorded information – as:

> To determine the exact class to which the book belongs, to place it where it can be found again at once when inquired for, to open to the scholar seeking information the entire resources of the library on a special theme, – these are constant duties of the librarian that imply a good system of classification.

And then, almost as an afterthought, Harris (1870, p. 114) added: "Every scheme of classification rests upon some philosophical system as its basis". Whereas most work on bibliographic classification theory has focused on the *how* of the first set of Harris' challenges, the exposition of the philosophical basis of classification have gone unnoticed for a good part of the last 140 years. It is not that work and research on bibliographic classification were not based on a solid philosophical foundation; but merely that it was assumed that the particular philosophical basis chosen did not constitute a philosophical basis. Miksa (1998) reviews the philosophical basis of significant bibliographic classification research in the twentieth-century and finds that several assumptions remain, including:

- Within most bibliographic classification research it is "assumed that knowledge categories are by nature hierarchical and logical in a classical, systematic sense" (Miksa, 1998, p. 81).
- Most bibliographic classification research subscribes "to the idea that somewhere, somehow, we can, or should try to, produce the one best classification system that will serve all purposes" (Miksa, 1998, p. 81).

These two assumptions are tightly linked into a common research paradigm; it is assumed that knowledge is neatly classified into logical categories, and that these categories are accessible and knowable in a fashion in which it is possible to replicate the structure in a single classificatory structure that can be used to organize all information for all purposes. The goal of this line of research has:

> [. . .] focused to a large degree on the creation of standard models, functioning as universal frameworks for classification, from which more local, for instance national or subject-oriented schemes could be developed (Albrechtsen, 2000, p. 2).

This line of work on bibliographic classification has, according to Svenonius (1992), historically searched for commonalities and generalities in two aspects:

(1) it has searched for commonalities across different domains; and

(2) it has searched for general laws and principles common to all classification systems.

These universal objectives have guided work in the area, and in many situations probably served as pragmatic goals at the expense of the "philosophical system" that Harris argued that all classifications rest upon. The library literature has generally not problematized the differences and commonalities of different domains and seems to have assumed that general laws and principles exist. Research on specialized libraries and information services has challenged the assumptions and begun working out frameworks that are domain specific and pluralistic; the work of the British Classification Research Group in the 1960s and 1970s (cf., e.g. Foskett, 1974; Vickery, 1975; Langridge, 1976) and Hjørland *et al.* recently (cf., e.g. Hjørland and Albrechtsen, 1995; Hjørland, 2002) are good examples of this line of work.

At this point, it might be necessary to provide an outline of issues dealt with in knowledge organization research and practice broadly speaking[2], I divide knowledge organization research and practice into three sets of knowledge organization problems (KOPs):

(1) *Big KOP*. This is the organization and representation of large quantities of information for unrecognizable many people; people with varied interests, beliefs, positions, knowledge, expertise, etc. The web is the prototypical example of such a KOP, large academic and many public libraries are also Big KOPs.

(2) *Medium KOP*. These are information collections for particular, stated, clear, objective, and specific purposes to be used by people with particular, similar interests, beliefs, positions, knowledge, expertise, etc. which can be known, understood, and articulated by those in charge of the collection or service. A company's intranet, a web portal, a store, and some special libraries are typical examples.

(3) *Small KOP*. This is individuals' information management challenges and collections. These KOPs are particular to an individual's (or a few individuals') personal information collection and will typically be managed by the same individual(s). The information could be e-mails, documents, files, photos, etc. which will be collected, searched, and used by individuals for their own usages.

The past 150 years of work on bibliographic knowledge organization has, more or less, focused on creating systems that could address the Big KOP with the development and research into universal classification systems and the work has been guided by the aforementioned assumption and goals outlined by Miksa and Svenonius. There has been some work done on Medium KOPs – especially with the development of special and domain-based systems and the development and research into techniques, methods and approaches to the design and development of specialized controlled vocabularies often in line with approaches outlined by the Classification Research Group and Hjørland *et al.*

We have seen some interests over the years in Small KOPs, personal information management, and it could seem as if this area is developing into a specialized sub-discipline in information science, with its own research agendas, etc. (Jones and Teevan, 2007).

Research and inquiries into Big and Medium KOPs have historically been merged into one issue, and treated as if the challenges are of equal nature. However, if it assumed that classifications grow out of people's activities and particular events and actions, then classification work that rests on such a conceptual foundation of classification is directed by the activities of people and as such demand the possibility of articulation and realization of those activities by designers and constructors of classifications. While it is possible to articulate and realize the activities that support Medium KOPs, it is not possible for Big KOPs. The two sets of challenges and issues, Big and Medium KOPs, are quite distinct and inquiries into them should be separated into two distinct areas of focus with distinct vocabulary, interests, and agendas. I see two main challenges with the current paradigm of merging Big and Medium KOPs:

(1) Universal systems and standards for Big KOPs were mainly developed with the assumptions for Medium KOPs in mind. The sort of KOP that the universal systems and standards were created to address was much smaller in scale 150 years ago, even 40 years ago, while the KOP grew in complexity – and has grown tremendously lately – the Big KOP is still addressed from the same epistemological approach and using the same systems and assumption as when the KOP was much smaller. There may be financial and practical reasons for

keep using the systems of yesterday[3] but social interaction applications, thinking, and approaches will eventually – and has to a large degree already – demonstrate other avenues to address the challenges and issues and suggest new approaches to make information accessible.

(2) Medium KOPs grows in complexity, interests, importance, and demand. There are good reasons to develop robust knowledge organization systems for information collections and services that are used in particular domains, for particular interests, by specific people and the demand will likely grow. More research is needed to increase our understanding of such Medium KOPs. While much research into knowledge organization systems has followed the Big KOP tradition and focused on the *systems*, Medium KOPs are characterized by uniqueness in domain, interests, and people and while there are certain characteristics of the systems that remain applicable regardless of context, much needs to be understood in terms of the relation between activities and knowledge organization.

This distinction between different kinds of KOPs is helpful to understand the various contexts that classification supports and it helps to recognize the characteristics of different sets of challenges. It also challenges, from a practical standpoint, the assumption that the one best system can be discovered, developed, and deployed. The distinction does not, however, provide a conceptual foundation for classification, for that to emerge we need first to step back and review the first principles of classification.

Folk theory of classification

Most work in library and information science classification research is based on a folk, or naïve, theory of classification that builds on a broader folk theory of categorization, which "says that things come in well-defined kinds, that the kinds are characterized by shared properties, and that there is one right taxonomy of the kinds" (Lakoff, 1987, p. 121); this folk theory of classification roots the aforementioned assumptions mentioned by Miksa (1998). Likewise, most bibliographic classification theory stipulates that documents are holders of concepts and concepts are context and human independent constructs[4] and that classification brings together concepts based on similarity. The notion of similarity as the defining notion runs deep in bibliographic classification and has been used since the dawn of modern bibliographic classification[5], through the formation of recent approaches[6], to contemporary work[7,8]. The notion of similarity, or likeness, has been criticized for being insufficient to capture the complexity of classification and as such, "likeness contributes little to an explanation of classification" (Broadfield, 1946, p. 3).

According to Broadfield, the basic challenge of using similarity as the defining notion for classification is that "likeness is not a quality of things. It is a relation between them, not a 'characteristic' of things" (Broadfield, 1946, p. 2). We attribute characteristics to things when classifying them and the classifications we create reflect a particular purpose or worldview, as Borges' Chinese classification of animals showed. Following this line of thought:

> [. . .] in certain conditions men may be like machines, or laughing-stocks, or obstacles, but in these cases the likeness, though unmistakable, does not reveal the kind of thing we think man to be (Broadfield, 1946, p. 5).

In other words, a particular classification is not given by the things classified, it is not enough to analyze the things for their essences and discover the true nature of things. Classifications are built by humans, for specific purposes, in specific contexts, and "the whole system is intuited, not hacked out characteristic by characteristic" (Broadfield, 1946, p. 5). The main challenge with using similarity as the main defining concept for classification is that any two objects can be similar in some aspects. A lawnmover and a plum, for instance, are similar in the sense that "they both weigh less than 10,000 kg, both did not exist 10 million years ago, both cannot hear, both can be dropped, both take up space and so on," which can only lead the conclusion that, "similarity is too flexible to explain categorization" (Bryant, 2000, p. 57).

To make similarity, the core defining concept of classification limits the power of the *classifier* to simply identifying *the* characteristics of things, and grouping those things that share characteristics. Classification thereby becomes a technical process that is merely occupied with the discovery of what things are and placing them in the one system given by nature. The idea that things possess properties that are objectively discoverable and that things in the world fall naturally into certain categories is a conception that has been with us for a long time:

> From the time of Aristotle to the later work of Wittgenstein, categories were thought to be well understood and unproblematic. They were assumed to be abstract containers, with things either inside or outside the category. Things were assumed to be in the same category if and only if they had certain properties in common. And the properties they had in common were taken as defining the category (Lakoff, 1987, p. 6).

This idea that the properties of things and categories to which things belong are given and universal forms the generic folk theory of classification, and feeds the idea that *likeness* is the defining, central notion in bibliographic classification work.

If, however, similarity is disbanded as the core concept in classification and:

> [. . .] we need to explain *why* we choose certain attributes over others on which to base our similarity judgements and to explain *in virtue of what* we judge members of a category to be similar. Why, for example, do we classify a Bedlington Terrier as a dog when it seems to have as many similarities with a lamb as with a Great Dane? (Bryant, 2000, p. 57).

These are serious questions that have profound impact on the practice of bibliographic classification; though I am less sure that the bibliographic classification research community has been able to satisfactorily address these questions and contribute to a solid theory and practice for classification. A solid theory and practice of classification should be able to address the basic question of, "why do we choose certain groups of similarity attributes over and above other for classification purposes?" (Bryant, 2000, p. 61). Since there is no one objectively correct way to classify a set of entities, classification systems should be based on an explanation for how the classes and relationships between classes have been established.

To enable a move toward explanation-based classification theory and practice, it is necessary to require and appreciate the inherent interpretative basis of classification work. As Eco (1984, p. 46, 1997) reminds us, "a sign is not only something that stands for something else; it is also something that can and must be interpreted," we must, similarly, remember that a class, and a classification, is not only something that stands for something; a class must and can be interpreted. Classification work is more than merely matching an item with a certain class and other items already in that class;

it requires an interpretation of the item, the class, and their relations. It requires that the classifier understands the class and the item, and its context. It requires that the classifier understands possible ways an item could be classified and is able to articulate why a particular classification was chosen. Hjørland (1997, p. 111) gives an example of this challenge:

> A stone on a field contain different information for different people (or from one situation to another). It is not possible for information systems to map *all* the stone's possible information for every individual. Nor is any *one* mapping the "true" mapping.

This line of thinking shifts classification from a task in which the characteristics of an item is identified and used as basis for classification to one in which one considers the contexts and usages of the items – and classifies the items based on their roles in discourses. Instead of asking, "what it is?" one asks, "what is it used for in this context?"

It is a common dictum among librarians that "value judgments have no place in the creation of bibliographic files" (Hagler, 1997, p. 43), it is assumed that the role of librarians is a *liberal* one, in which:

> [. . .] the librarian not only has no politics, no religion, and no morals; he has no opinion on any open question. Librarians see their role as one of complete hospitality to all opinions (Wilson, 1983, p. 190).

But classifications *are* epistemological statements; they *do* say something about the world, and they embed politics, religion, and moral. Weinberger (2007, p. 56) asks us to imagine what would happen if the editors of the Dewey system finally decided to do something about its biases:

> Imagine that the system's editors decide to fix the system once for all. They move discredited categories such as phrenology to the right of the decimal place. They consolidate the Christian topics, pull Buddhism up a couple of integers, push Baha'i down, drag computer science into technology, demote philosophy from its top-ten status because, frankly, philosophy isn't the queen of the sciences anymore, and do a thorough housecleaning. What happens next?

It might be important here to remember that it is:

> [. . .] impossible to understand the placement of topics in DDC and LCC if one did not know that both are organized around nineteenth century views of the academic disciplines. In the last century, psychology was considered in universities be a subdivision of philosophy and thus not an empirical testable field of study. Consequently, psychology appears in the DDC 100s [. . .] psychology seems solely misplaced one hundred years after the [system was] invented (Beghtol, 1997, pp. 92-3).

It is clear that keeping up to date with the scholarly developments and scientific revolutions requires constant evolvements of a classification system and in some cases, as Weinberger and Beghtol point out, significant rearrangements of the structure. However, regardless of the outdated view of the academic disciplines, Weinberger (2007, p. 56) continues and explains what might happen next:

> Tens of thousands of librarians around the world pick up their razor blades and scrape the white numbers off the spines of millions of books, muttering under their breath about those damn editors who don't understand that every little change means that librarians inhale toxic

> white dust. Entire card catalogs get discarded, so to speak, and millions of new cards printed up. Books are piled up, moved from this shelf to that. And at the end of months or years of work, the complaints begin. The Sunnis and Shiites are upset because they've been put at the same level. The Jews are furious because the Jews for Jesus, whom they view as Christian predators, are listed under Judaism. Feminism and fundamentalist Christians find themselves making common cause to get studies of pornography removed from the arts section. East Somewhere is furious because it doesn't recognize West Somewhere as a legitimate country. Librarians are out buying razor blades in bulk and white ink by the gallon.

It is not possible to "fix" the system. If by fixing we mean: creating a system that is neutral and holds no values.

It is not a challenge or a fault unique to the Dewey system, it is not a problem that is unique to library classification; it is a common challenge for all classification systems. No classification system is objective, value neutral, and a representation of how things really are, as Borges reminded us. Any classification is a classification from a particular point of view, for a particular purpose. It is often noted that, "classifications can only be judged in terms of the objectives of those who utilize them, and are thus best constructed with careful attention to those objectives" (Szostak, 2008, p. 320).

The objective of a classification needs to be articulated; it is not possible to judge the Dewey system without a clear understanding of the system's premise, point of view, history, and objectives. Further, objectives are always someone's objectives. An objective of, say, organize the world's knowledge, will always be an objective of someone, somewhere, at some time. These constraints should be articulated and explained. One cannot assume that his or her objective to organize the world's knowledge is neutral and value free, as Hagler implies librarians should.

The distinction between a positivistic and a pragmatic approach to classification, which has gained some attention[9] lately, is a false distinction, just like the distinction Bliss makes between a natural, a developmental, and a pedagogical order was unfruitful (Bliss, 1929). The real question, the one that will make a different for the success and acceptance of a classification is whether it is *trusted*. The Linnaean classification is a trusted classification, the periodic table is a trusted classification, just like the Women's Thesaurus is trusted. These systems are not trusted because they represent the world better, more correct, or more fully than other systems, but because there is a clear connection between the system's philosophical basis – as Harris asked for – and the system's structure. These systems are treated and thought of as statements to open questions; it is assumed, and argued, that there are alternative ways such systems could be structured. The Dewey system is being ridiculed not just because it a biased system or because it is a bad system by some measure, but because it assumes a neutrality where there is none. It could seem as if the Dewey system treats its structure as a reaction to a closed question; that there is one best way to structure the classification. In explaining its effort to keep the system updated, the Dewey (2009) system explains:

> Keeping pace with knowledge and helping you organize it is what the DDC system is all about. Our editorial staff and international policy committee work closely with librarians to ensure the revisions, expansions and additions we make to DDC truly reflect the needs of today's library users. It's input from experts like you that continues to make Dewey timely, relevant and indispensable.

It is comforting to know that the editors of Dewey collaborates with a number of people in identifying "the needs of today's library users," however, the basic assumption

behind the statement is problematic. The statement assumes that there is a specific set of needs that "today's library users" have, and that it is possible for the editorial team and the international committee to understand and plan according to these needs. The editors seem to act as if the Dewey system exist to address some Medium KOP somewhere, for some body. Some librarians have, in fact, started dispensing Dewey with alternative methods of organizing their collections. A growing number of libraries have begun to use the Book Industry Standards and Communications System or simply by sorting books "from book drops right onto trucks" (Fister, 2009) that are placed near the library's entrance, these books never make it back to the stacks and into Dewey's categories (Fister, 2009).

Dewey's statement that the editorial team works with a number of "experts" to ensure that the system "truly reflect" the needs of library users is an attempt to establish the credibility and trustworthiness of the system. It is unclear what the experts are experts at and what claims they have to expertise in this area – do they have expertise in librarianship, or are they also experts in the subject matter to be classified? I will return to the issues of experts and expertise later in the paper, but it is important here to note that without this credibility and trustworthiness more librarians might dispense the system and move to alternative approaches. Although the Dewey system does call for "input," it should be noted that unlike the Linnaean system, the periodic table, the Women's Thesaurus, or LibraryThing for that matter, the Dewey system does not delegate the authority to name the categories, create the relationships, and classify the material to the community. It is assumed that the editors and the policy committee along with librarians have the authority to create a trustworthy classification.

As bibliographic classification moves away from the classic, folk theory and towards a framework for classification where the social sphere and its interactions, conversations, interests, and discourses become the focal point, and where it is accepted that the particular structure of the classification is just one among many possible, it is becomes important to rethink how to trust classification and how to design authoritative classification systems.

Trustworthy classifications

Designers and editors of classification systems have an *administrative authority* qua their positions. The editorial team at Dewey possesses an administrative authority that is given by their positions and the role the Dewey system plays. Some body has chosen the editors to design and edit the system; to ensure that the system is good, successful, adheres to the standards, and is within budget. Designers and editors have "a recognized right to command others, within certain prescribed limits" (Wilson, 1983, p. 14) and we will trust these administrative authorities with respect to certain technical tasks. However, designing and editing classifications are not merely a technical task; it is a task that involves making ontological statements about the world and the relations among entities of the world. While administrative authorities can be trusted with the technical aspects of designing and editing classifications, only *cognitive authorities* can be trusted to make ontological statements.

Cognitive authorities are those people we turn to for knowledge, insight, and advice on particular matters. These people have gained their authority not by being chosen by some body, but because we chose to let them influence our thinking. We recognize that there are certain people whom we trust on particular matters and other people we trust

on other matters. It is seldom that we trust the same people on all matters. Wilson (1983, p. 18) defined a cognitive authority as:

> One to whom we turn for information but also one to whom we turn for advice, even (or particularly) in cases where it is clear that there is no knowledge to be had at all. Cognitive authorities are valued not just for their stocks of knowledge (answers to closed questions) but for the opinions (answers to open questions) and for the advice on the proper attitude or stance on questions and their proposed answers. Cognitive authority is not limited to the provision of knowledge or information, stopping when the limits of available knowledge are reached. Cognitive authority can extend to any sort of question: moral, religious, political, aesthetic, technical, scientific, philosophical – and can be exercised in areas where all questions are open and expected to remain open indefinitely.

Many, if not all, classifications make ontological statements that someone, somewhere, at some point will object to. They are responses to open questions. As Weinberger's tale about librarians rushing out to purchase razor blades points out, it is impossible to create classifications are neutral, unbiased and which will satisfy everyone, everywhere. Designers and editors must consult with cognitive authorities within the classification's domain and objectives to develop explanations for why the classification is structured as it is, why the classes have the names they have, and articulating the system's philosophical basis which Harris (1870) reminded us that all classification systems rest upon.

While we must recognize that expertise and experts do play an important role in creating and maintaining classification systems, the question raised here is on which basis they can do so and claim trustworthiness. While the Dewey system might seek input from "experts like you," we need to remember that being an expert is different than being a cognitive authority for "one can be an expert even though no one else realizes or recognize that one is" (Wilson, 1983, p. 13), when we turn to particular people for information and advice, we turn to those we *recognize* as knowledgeable and trustworthy. In most cases, these people are experts within a given area, but we turn to them because we trust them, not because they are experts. Lakoff (1987, p. 123) points out, "There is some body of people in society who have the right to stipulate what words should designate, relative so to some domain or expertise," and while we do grant *administrative authority* to editorial teams of classification systems, the decisions about what categories mean, their structure, their relationships and what goes into various categories must be done by *cognitive authorities* whom users trust. The structure of classification systems and their categories are often the result of broad consultation and open feedback mechanisms that have provided editorial teams with data about the systems and their users, as Bowker and Star (1999, p. 44) note:

> [...] what appears as universal or indeed standard, is the result of negotiations, organizational processes and conflict. How do these negotiations take place? Who determines the final outcome in preparing a formal classification? [...] Someone, somewhere, must decide and argue over the minutiae of classifying and standardizing.

These negotiations of the minutiae of classification systems are often lost, when classification systems are presented as finished, polished, and final products, as neutral, objective epistemic statements about the world as it is. Again, the question remains, what make classification systems trustworthy?

A number of scholars have recently discussed ways to make library and information systems trustworthy. Lankes (2008) argues that credibility should be extended to consist of two interdependent aspects, authority and reliability. He finds that credibility is established through multiple sources and as such in a "conversation [that] is open and ongoing" (Lankes, 2008, p. 681). The consequence of this is that systems should be designed to facilitate participation and open dialogue. Kelton *et al.* (2008, p. 371) argue that the focus of trust in information needs to change from the qualities of information that make it trustworthy to the understanding the "perceptions of the person who is using that information". Rieh (2002) found that people do take cognitive authority into account when making judgments about the quality of information found of the web and Hertzum *et al.* (2002) found that trust plays a significant role when people access and choose among various information sources. van House (2002) states that credibility and trust are bound by epistemic communities and established within these communities, thus one must become a member of the particular community to appreciate and determine which sources are trustworthy.

However, as noted by Fallis (2008, p. 1665), "library and information scientists often focus on the *descriptive* question of when people actually do grant *cognitive authority* to information sources" and similar to Fallis (2008, p. 1665), the concern here is "with the *normative* question [. . .] of whether people *ought* to grant cognitive authority" in Fallis' case to wikipedia, here to a classification system. The focus here is to establish ways in which classification systems can be designed so that people ought to regard them with trust. The concern is less with how and when people trust systems; while this descriptive aspect is of interest, it more relevant when people attempt to verify a piece of information and of less interest when establishing or designing a system. Though descriptive aspects will, of course, influence the design approach.

Before we move on and consider the way classification systems can be designed to exhibit trust, it may be useful to consider ways in which people can or should grant cognitive authority to classification systems and trust them. A number of scholars (cf., e.g. Fallis, 2004, 2008; Goldman, 1999, 2001; Hardwig, 1991) have discussed the concept of trust in detail and the epistemology of testimony approach proves especially helpful here. Based on this literature, I would think that we could determine ways in which people can find out whether to trust a particular classification. For instance, people could evaluate the trustworthiness of a particular classification by asking questions in five areas:

(1) *Credentials*. Does the people or organization behind the system hold positions and academic degrees that entitle them to be trusted?

(2) *Record*. Has the people or the organization behind the system previously published reputable classifications?

(3) *Reputation*. Does the people or organization behind the system have a reputation for producing quality and trustworthy classifications?

(4) *Witness*. Does other people use this particular classification?

(5) *Agreement*. Do other classification systems have similar classificatory structures?

Experts in a particular domain might be able to apply their knowledge and expertise of the domain in considering whether they trust a particular classification. They will not

necessarily look for whether the classification is correct or align with their own view of the domain, but rather whether the classification make sense, whether the classification is well reasoned, well explained. They can make this judgment based on implicit or explicit *esoteric* statements made about the classification system in the epistemic community in relation to the five areas mentioned above, or any other criteria they chose to judge the classification by.

Laypersons, or people with little knowledge of the particular domain, on the other hand, must rely on indirect information when addressing the five questions about the trustworthiness of a classification system because the esoteric statements are inaccessible to them. For laypersons to assess the trustworthiness of a classification system, it needs to be possible for them to assess the five areas through *exoteric* statements (Goldman, 2001). In other words, they need to able to evaluate the classification system's trustworthiness in a language that is "outside the domain of expertise" (Goldman, 2001, p. 94). Experts can assess the trustworthiness through inspection of the systems combined with the discourse in epistemic community and thereby gain enough understanding and knowledge to address the five areas of questions directly. Laypersons, on the other hand, do not have this advantage and they need to access information about the five areas indirectly. In other words, they need to trust someone to give them access to information that they can use to evaluate the trustworthiness of a classification.

For the verification of information and testimony in general, laypersons might be able to compare and contrast various sources; they can consult multiple sources of information and assess the trustworthiness of the information by verifying that several sources provides the same answer to a question, for instance (Fallis, 2004). When assessing the trustworthiness of classification systems, people are usually not in a situation where they can choose between several classifications, they are offered a single classification that they need to use to retrieve information, to make sense of some particular phenomena, or to understand a particular domain. It is therefore important that the system provides users the ability to assess its trustworthiness in *exoteric* terms.

We can now return to the question of how to design trustworthiness classification systems.

Previously, we saw that a good deal of effort in thinking about classification has gone into the debate about whether a classification can or should represent things as they really are. We saw that there has been a debate about whether classification should strive towards being neutral and value free, what to do with bias in classification, and whether bias ought to be embraced and accepted as a foundation for all classifications. And we have seen how explanations are required for people to understand classification systems when we take pluralism for granted. That led to the conclusion that classification should not strive towards being correct, but towards being trustworthy. For people to trust a classification system, they need to have information about the decisions, principles and philosophy that have informed the classification, as well as information about the people and organizations behind the system. This leads me to formulate the *principle of transparency* for classification systems.

The principle of transparency says any classification should make available to its users statements about the basis on which the system is designed. Designers and editors need to appreciate the plurality of the epistemic and ontological statements their classifications make, and should guide the users through their negotiation and decisions. Designers and editors should bring their classification out of their contexts

to reveal their nature, and they should make the foundation of the classifications available to both experts (through esoteric statements) and to laypersons (through exoteric statements). Only classifications that are transparency can be trusted. Notice that the challenge for designers and editors are not to correct their classifications, or ensure that they are in line with contemporary theory and thought, but simply to provide explanations.

People ought to grant cognitive authority to classification systems that are transparent. People do not have to agree with the structure of the classification to trust the system; they could use Borges' classification of animals or DDC's religion section if they understand and appreciate the reasons behind the structure and the categories. As Lankes (2008) noted, an important way to create credibility is by being open and establishing a dialogue. Systems that are closed and appeal to their administrative authority to gain cognitive authority fail and becomes less trustworthy. Those that acknowledge their "arbitrary and conjectural" (Borges, 1952, p. 104) nature and explain it, becomes more trustworthy.

Conclusion

In a world where knowledge and truth are social constructions and where these notions are thought of, produced by and established in social interactions, classifications need to be developed in constant collaboration with its users. Classifications are part of already of existing conversations and activities; their trustworthiness depend on their ability to share the conceptual basis with these conversations and activities.

As bibliographic classification work embrace the notion of plurality and accepts that all classifications are biased in some way, there will be an increased strive towards establishing the trustworthiness of classification systems. The fact that we accept a pluralistic foundation, should not mean that that an "everything goes" approach is accepted as well. While users previously might have accepted some administrative authorities as cognitive authorities, it can be expected that authorities will be challenged even further in the future and that classification systems, along with other authorities in society, need to established their trustworthiness in new ways; they need to be transparent.

While the conceptual foundational for all classification needs upon the principle of transparency, the exact way that this plays out depends largely on the context, expectations, domain, users, and the type of KOPs that the classification addresses. While neutrality and non-bias cannot be expected of any classification, it is possible for designers and editors to grasp, understand, and appreciate the problem sphere for Medium KOPs and they can strive towards systems that respect and to some degree reflect the domain's structure (Hjørland and Albrechtsen, 1995). Designers and editors of Big KOP systems, on the other hand, have an increased duty to handle their systems' inherent biases responsible (Feinberg, 2007). They must adhere to the principle of transparency to create systems that people will trust and must to a greater extend than Medium KOPs involve their users directly in the development of the classification structures.

Classification theory has moved award from ontological discussions of whether classification systems ought to reflect things as they really are and have moved towards a relativistic foundation. It has embraced the fact that all systems are inherently biased. To regain trust and authority designers and editors of classification

systems must embrace the principle of transparency and explain their decisions and show the conceptual and philosophical foundations for their systems.

Notes

1. See Scerri (2007) for a detailed analysis and discussion of the Periodic Table. There are clearly many problematic issues with the Periodic Table.
2. I use the term "knowledge organization" to denote the broader sphere of issues dealt with in the representation and organization of information and other artifacts for retrieval. I use the term "classification" more narrowly to denote the acts of establishing the sense of various entry points and their relationships.
3. To borrow Albrechtsen's (2000) phrase.
4. The ISO Standard (1985) on subject representation, for instance, defines *subject* as: "Any concept or combination of concepts representing a theme in a document" and *concept* as: "A unit of thought."
5. "The 'putting together of like things' is [...] the fullest and most exact' definition of classification" (Richardson, 1935, p. 1).
6. "Classification [...] means putting together things or ideas that are alike" (Vickery, 1975, p. 1).
7. "Classification brings like things together" (Svenonius, 2000, p. 10).
8. Other related definitions include: "Classification is the act of bringing like things together" (Buchanan, 1979, p. 9); "The purpose of classification is to bring related items together" (Wynar, 1992, p. 317); "To group all works of a kind together" (Chan, 1994, p. 259).
9. Szostak's (2008) polemic paper which discusses a paper by Hjørland and Nissen Pedersen (2005) – and Hjørland's (2008) reply to Szostak are good examples of the attention this has been given.

References

Albrechtsen, H. (2000), "Who wants yesterday's classifications? Information science perspectives on classification schemes in common information spaces", *Proceedings of the International Workshop on Cooperative Organization of Common Information Spaces, Copenhagen, Denmark, August 23-25*.

Beghtol, C. (1986), "Semantic validity: concepts of warrant in bibliographic classification systems", *Library Resources & Technical Services*, Vol. 30 No. 2, pp. 109-25.

Beghtol, C. (1997), "'Itself an education': classification systems, theory, and research in the information studies curriculum", *Technical Services Quarterly*, Vol. 15 Nos 1/2, pp. 89-107.

Bliss, H.E. (1929), *The Organization of Knowledge and the System of the Sciences*, Henry Holt, New York, NY.

Borges, J.L. (1952), "The analytical language of John Wilkins", *Other Inquisitions 1937-1952*, Souvenir Press, London, 1973.

Bowker, G.C. and Star, S.L. (1999), *Sorting Things Out: Classification and Its Consequences*, MIT Press, Boston, MA.

Broadfield, A. (1946), *The Philosophy of Classification*, Grafton, London.

Bryant, R. (2000), *Discovery and Decision: Exploring the Metaphysics and Epistemology of Scientific Classification*, Associated University Presses, Cranbury, NJ.

Buchanan, B. (1979), *Theory of Library Classification*, Clive Bingley, London.

Chan, L.M. (1994), *Cataloging and Classification: An Introduction*, 2nd ed., McGraw-Hill, New York, NY.

Dewey (2009), "Dewey services", available at: www.oclc.org/dewey/overview/default.htm (accessed October 31, 2009).

Eco, U. (1984), *Semiotics and the Philosophy of Language*, Macmillan, London.

Eco, U. (1997), *Kant and the Platypus: Essays on Language and Cognition*, Secker & Warburg, London.

Fallis, D. (2004), "On verifying the accuracy of information: philosophical perspectives", *Library Trends*, Vol. 52 No. 3, pp. 463-87.

Fallis, D. (2008), "Towards an epistemology of Wikipedia", *Journal of the American Society for Information Science and Technology*, Vol. 59 No. 10, pp. 1662-74.

Feinberg, M. (2007), "Hidden bias to responsible bias: an approach to information systems based on Haraway's situated knowledges", *Information Research*, Vol. 12 No. 4, available at: http://InformationR.net/ir/12-4/colis/colis07.html (accessed November 1, 2009).

Fister, B. (2009), "The Dewey Dilemma: in the search for better browsability, librarians are putting Dewey in a different class", *Library Journal*, October 1, available at: www.libraryjournal.com/article/ca6698264.html (accessed November 1, 2009).

Foskett, D.J. (1974), *Classification and Indexing in the Social Sciences*, Butterworths, London.

Goldman, A.I. (1999), *Knowledge in a Social World*, Oxford University Press, Oxford.

Goldman, A.I. (2001), "Experts: which ones should you trust?", *Philosophy and Phenomenological Research*, Vol. 63 No. 1, pp. 85-110.

Hagler, R. (1997), *The Bibliographic Record and Information Technology*, ALA, Chicago, IL.

Hardwig, J. (1991), "The role of trust in knowledge", *The Journal of Philosophy*, Vol. 88 No. 12, pp. 693-708.

Harris, W.T. (1870), "Book classification", *Journal of Speculative Philosophy*, Vol. 4, pp. 114-29.

Hertzum, M., Andersen, H.H.K., Andersen, V. and Hansen, C.B. (2002), "Trust in information sources: seeking information from people, documents, and virtual agents", *Interacting with Computers*, Vol. 14 No. 5, pp. 575-99.

Hjørland, B. (1997), *Information Seeking and Subject Representation: An Activity-Theoretical Approach to Information Science*, Greenwood, Westport, CT.

Hjørland, B. (2002), "Domain analysis in information science: eleven approaches – traditional as well as innovative", *Journal of Documentation*, Vol. 58 No. 4, pp. 422-62.

Hjørland, B. (2008), "Core classification theory: a reply to Szostak", *Journal of Documentation*, Vol. 64 No. 3, pp. 333-42.

Hjørland, B. and Albrechtsen, H. (1995), "Toward a new horizon in information science: domain-analysis", *Journal of the American Society for Information Science*, Vol. 46, pp. 400-25.

Hjørland, B. and Nissen Pedersen, K. (2005), "A substantive theory of classification for information retrieval", *Journal of Documentation*, Vol. 61 No. 5, pp. 582-95.

ISO Standard (1985), *Documentation – Methods for Examining Documents, Determining Their Subjects and Selecting Indexing Terms*, International Organization for Standardization, Geneva, 5963.

Jones, W. and Teevan, J. (2007), *Personal Information Management*, University of Washington Press, Seattle, WA.

Kelton, K., Fleischmann, K.R. and Wallace, W.A. (2008), "Trust in digital information", *Journal of the American Society for Information Science and Technology*, Vol. 59 No. 3, pp. 363-74.

Lakoff, G. (1987), *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, University of Chicago Press, Chicago, IL.

Langridge, D.W. (1976), *Classification and Indexing in the Humanities*, Butterworths, London.

Lankes, R.D. (2008), "Credibility on the internet: shifting from authority to reliability", *Journal of Documentation*, Vol. 64 No. 5, pp. 667-87.

Miksa, F. (1998), *The DDC, the Universe of Knowledge, and the Post-Modern Library*, Forest Press, Albany, NY.

Richardson, E.C. (1935), *Classification: Theoretical and Practical*, H.W. Wilson Co., New York, NY.

Rieh, S.Y. (2002), "Judgment of information quality and cognitive authority in the web", *Journal of the American Society for Information Science and Technology*, Vol. 53 No. 2, pp. 145-61.

Scerri, E.R. (2007), *The Periodic Table: Its Story and Its Significance*, Oxford University Press, Oxford.

Shirky, C. (2005), "Ontology is overrated: categories, links, and tags", *Clay Shirky's Writings about the Internet. Economics & Culture, Media & Community*, available at: www.shirky.com/writings/ontology_overrated.html html (accessed November 1, 2009).

Shirky, C. (2008), *Here Comes Everyone: The Power of Organizing without Organizations*, Penguin, New York, NY.

Svenonius, E. (1992), "Classification: prospects, problems and possibilities", *Classification Research for Knowledge Representation and Organization*, Elsevier, New York, NY.

Svenonius, E. (2000), *The Intellectual Foundation of Information Organization*, MIT Press, Cambridge, MA.

Szostak, R. (2008), "Classification, interdisciplinarity, and the study of science", *Journal of Documentation*, Vol. 64 No. 3, pp. 319-32.

van House, N. (2002), "Digital libraries and practices of trust: networked biodiversity information", *Social Epistemology*, Vol. 16 No. 1, pp. 99-114.

Vickery, B.C. (1975), *Classification and Indexing in Science*, Butterworths, London.

Weinberger, D. (2007), *Everything is Miscellaneous: The Power of the New Digital Disorder*, Times Books, New York, NY.

Wilson, P. (1983), *Second-hand Knowledge: An Inquiry into Cognitive Authority*, Greenwood, Westport, CT.

Wynar, B.S. (1992), *Introduction to Cataloguing and Classification*, 8th ed., Libraries Unlimited, Littleton, CO.

Corresponding author
Jens-Erik Mai can be contacted at: je.mai@utoronto.ca