Rick Szostak, University of Alberta

Classifying for Diversity

Abstract

This paper argues that a new approach to classification best supports and celebrates social diversity. It maintains that we should want a classification that both facilitates within-group communication and cross-group communication. This is best accomplished through a truly universal classification that classifies works in terms of authorial perspective. Strategies for classifying perspective are discussed. The paper then addresses issues of classification structure. It follows a feminist approach to classification, and shows how a web-of-relations approach can be instantiated in a classification. Finally the paper turns to classificatory process. The key argument here is that much (perhaps all) of the concern regarding the possibility that classes can be subdivided into subclasses in multiple ways (each favored by different groups or individuals) simply vanishes within a web-of-relations approach. The reason is that most of these supposed ways of subdividing a class are in fact ways of subdividing different relationships among classes.

Introduction

Classification necessarily involves structure, and thus there is a perhaps inevitable tension between classification and the celebration[1] of diversity. Since not all societal groups – gender, ethnic, or socio-economic – have been equally represented in the development of major classifications, a related concern arises that existing classifications privilege certain ways of looking at the world while obscuring others.

Can we develop a universal classification that reflects and celebrates diversity? This question is rarely asked explicitly, despite its obvious importance. It would be unfortunate to assume a needlessly pessimistic answer to this question. It could be that some of the advocacy of the exclusive performance of domain analysis in the literature reflects an unstated view that each societal group is best served by its own classification.[2]

This paper will first explore the purpose of classification from the perspective of diversity: what exactly should we want a classification to do in order to respect and support diversity? Once we have set goals for classification, we can then proceed to examine questions of structure: how should a classification be organized in order to achieve these goals? Finally we can turn to process: Is it feasible to develop such a classification?

Purpose

How can library classification best serve the interests of diverse groups? The first question that might be asked is whether we wish to encourage cross-group understanding, or is within-group understanding of paramount importance? If it was thought that each societal group should have unique and privileged access to its own literature, then domain analysis would be the obvious way to go. Each group could classify its own literature in terms of concepts with which non-group members would

---

[1] In a salute to ambiguity I appreciate that this is not quite the right word. I mean by it to recognize, appreciate, and support diversity to an appropriate degree.

[2] As I noted in Szostak (2013c), an argument that domain analysis is all that we can do should be carefully distinguished from an argument that it is all that we should want. They are each misguided for quite different reasons. I was inspired there and in this paper by Fox (2012).

be unfamiliar – or better yet apply unique meanings to terms which others might mistakenly believe to be familiar. Outsiders would then find it extremely difficult to access the literature. Only those within the group, and possessing a clear apprehension of the true meaning of the concepts employed in the classification (grounded in turn in how the concepts are employed in the group's literature) would be able to readily navigate the literature. Classification would support group solidarity.

The cost would be in terms of cross-group understanding. Group members would have to master other classifications if they wished to read in the literatures generated by other groups. And they would have to publish in venues classified in other ways if they wished to speak beyond their group. Information scientists could generate translation devices, perhaps, but even these would be costly to master for each group one wished to engage.

Cross-group understanding would be best facilitated by a universal classification. Then members of any group would have equal access to all literatures. The cost – we are leaving aside issues of practicality for the moment – would be that group members would lack any special access to the literature of their own group: works authored by group members might be hard to distinguish from works on similar topics by others. Especially if the group were small they might legitimately feel that a universal classification militates against a sense of common cause and identity.

The challenge for information science is that the answer to our opening question in this section is almost certainly "Both." Only the most xenophobic would wish to cut their group off completely from interaction with others (especially as individuals in the contemporary world increasingly have several cross-cutting group memberships). And even the most universalist in outlook can appreciate the value of people being able to readily communicate to others with certain shared characteristics.

The lesson for information science is clear but rarely stated: we should strive to facilitate both across-group and within-group communication. And this result holds for any type of group: disciplinary, gender, ethnic, occupational/class, religious, sexual orientation, and so on.

Is it possible to pursue both goals simultaneously? I have argued elsewhere (Szostak 2010) that domain analysis and the pursuit of a truly universal classification can be complementary approaches. The key argument is that complex concepts – those that are understood differently across groups – can be broken into basic concepts (this is where domain analysis can be employed) that lend themselves to broadly similar understandings across groups (Szostak 2011). Basic concepts generally refer either to the things we study (including theories and methods) or relationships among these. Works and ideas are then classified in terms of combinations of these basic concepts. Notably, a classification grounded in such basic concepts will not only aid users in finding works written by members of any group but will aid them in understanding those works (by translating complex concepts into basic concepts) (Szostak 2013a). But these arguments – though showing that domain analysis is far from incompatible with universal classification, and that a classification can support cross-group understanding – do not on their own provide an answer to our present question. There is a piece missing: can we signal within such a universal classification the group membership of authors?

This question in turn raises both philosophical and practical questions. Philosophically, authors may often not wish to be seen as speaking as a member of a particular group.[3] And so we need to appreciate that attempts to classify works in terms of group membership have a potential downside. Authors may be striving to generate universal understanding. More pragmatically they may worry that group identification will blunt their ability to reach out to others. Allowing authors themselves to decide whether their works should be given any particular group identification provides an imperfect solution (and only for the living).

On the practical level, it would seem that the classificationist needs to develop a list of group descriptors. But perhaps not: if the classification relies on combinations of basic concepts, then every possible group will already of necessity be classified in order to describe works about that group. Group membership could then be signaled by linking "from the perspective of" and a particular group.

There may, though, be an even better way to proceed. Many scholars have argued that (at least some) works should ideally be classified in terms of authorial perspective.[4] None that I am aware of have suggested a detailed classification of perspectives. Yet one clear implication of this literature is that group membership will often not be the best signal of perspective. Describing a work as 'gender studies,' or alternatively as applying 'feminist theory,' will send a more valuable signal than merely noting the gender of the author.

I have long urged the classification of works in terms of theories – and methods – applied, and this was a key component of the León Manifesto (2007). I confess that I have only belatedly emphasized classification in terms of discipline or interdisciplinary field of author. And I now appreciate that it is important to embrace a much wider set of perspectives.

A variety of dimensions might be useful in capturing the perspective of an author: inter/disciplinary; theoretical, methodological, rhetorical,[5] epistemological, ideological, aesthetic, ethical.[6] More broadly we wish with authorial perspective to capture key motives and beliefs of the author. A key challenge for the information scientist is that there are imperfect correlations across these dimensions: not all who apply feminist theory are female, and feminist theory is applied outside of gender studies. A work classified along only one dimension will be missed by users searching along another.

In sum it is likely possible to enhance both across-group and within-group communication through a universal classification that classifies works in terms of authorial perspective. But our ability to achieve both depends on our developing a useful classification of authorial perspective.

---

[3] We can set aside for the moment a concern voiced by some scholars that the pretence of a universalist perspective is itself a marker of male dominance: the principle of literary warrant would still guide us to recognize such an aspiration.

[4] Otlet, Mills, and Vickery are among those who have over the years urged classification of works in terms of perspective.

[5] Feinberg (2011) speaks in particular of logical argument (manipulation of evidence), ethos (incorporation of audience beliefs and values to establish trust), genre adaptation (adjustment of formal elements). Her purpose is to argue for an explicit authorial voice in classifications, but her argument can be used for classifying works by authorial voice.

[6] Clavier and Paganelli (2012) argue that we should classify works by stance: criticism, agreement, consensus, etc.

As a segue to our next section, it is useful to engage with arguments made by Mai (2011). He argues that contemporary approaches to classification (grounded in ontology) reflect a modernist view that imagines in a realist fashion that the things we study exist separately from those who study them. He instead recommends an epistemological approach to classification that appreciates subjectivity. Though I am epistemologically more confident than Mai that consensus is possible due to our ability to fairly accurately apprehend reality, I can nevertheless appreciate that Mai provides a further justification for classifying works by perspective: this will help to identify the biases that an author brings to the work.[7] But Mai is not sure what a classification grounded in subjectivity would look like (nor is Hjørland 2012). He might thus be skeptical of our ability to classify perspectives in a way that respects all perspectives.

More generally, Mai doubts that there can be consensus on the classes within any classification. Though I am again much more optimistic than Mai, it will prove useful to try to meet this concern as much as possible. That is, if different people or groups will (even just sometimes) disagree over the nature of classes, then we should strive to minimize the scope for disagreement. This we will do in each of the next two sections by simply limiting the degree of hierarchical organization, and focusing on the classification of basic concepts. Mai also urges transparency: it should be clear how a classification was developed (so that the user can evaluate whether/what biases drove its development). This principle will also guide us.

Structure

We have identified above some characteristics to be sought in a classification. But what sort of structure will best achieve these goals? We start this section by responding to a feminist critique of classificatory practice, and then show that – not surprisingly – the classification that responds to this critique serves the various goals outlined above.

Olson (2007) has suggested that hierarchy is more reflective of a 'masculine' perspective, and that a classification that blended hierarchy with a 'web of relations' approach would be more gender-neutral. Women, she argues, are more likely to see the world in terms of a web of relations. But Olson also argues that underprivileged social groups would likely also benefit from a less hierarchical approach to classification.

Is such a classification possible? This critical question has not been addressed in the detail that it deserves. It will be argued that such a classification is indeed possible, with reference being made to the Basic Concepts Classification that I have been developing (and was briefly outlined above; see Szostak 2013b). The key again lies in classifying works in terms of combinations of basic concepts. One work might be classified in terms of how a phenomenon A influences a phenomenon B in a particular manner Z. Such a work will be found easily by anyone interested in how A might affect B in manner Z, regardless of the user's group membership. It can also be found easily by anyone studying how B influences C who then becomes curious about how to

---

[7]   But it will not capture all sorts of bias. In Szostak (2004, ch.5) I attempted a classification of the various sorts of bias that affect scholars. In addition to those sorts of bias grounded in socio-cultural understandings are biases reflective of human cognitive and perceptual limitations, biases rooted in the need to participate in an ongoing conversation, biases reflective of the incentives authors face, and biases reflecting institutional structures.

encourage changes in B. And someone interested in how F influences G in manner Z might become interested in other cases of influence of type Z. So this sort of compound classification utilizing basic concepts in fact instantiates a web of relationships at the level both of works and of the key arguments expressed in works. A user can thus follow, if they wish, a complex set of issues from one work to another. As Olson notes, present classifications facilitate browsing only within a hierarchy; the proposed structure also facilitates browsing across hierarchies.

Hierarchy is still necessary in such a classification, but to a much lesser extent. Types of influence can be captured through combinations of some 100 basic types of influence that can be organized in just two levels of hierarchy (occasionally three; see Szostak 2012). The things we perceive can, at least in the human sciences, be captured in very compact schedules (Szostak 2011, 2013b). Natural science requires much more detailed hierarchies of species and chemical compounds.

Olson is concerned that hierarchical approaches tend to privilege "being a Y" over "not being a Y" in general, and "being male" over "not being male" in particular. A classification that did not distinguish males from females in any class except gender itself could obviate this concern. In place of present practice, in which male nurses and female engineers are treated as some sort of anomaly, the classification here would use linked notation: (nurse)(male) and (nurse)(female) would be classificatorily equivalent – as indeed would be (nurse)(transgendered).

Olson worries that hierarchy privileges deduction over induction. I have long argued that the best approach to classification blends induction and deduction. And indeed this is one key reason for urging us to blend universal and domain-analytical approaches (see Szostak 2010), for a universal classification demands some logical structure whereas domain analysis is inherently inductive. I urge a deductive approach to hierarchy below. Allowing elements in any hierarchy to be freely linked with elements in any other hierarchy provides immense scope for an inductive appreciation of any connection drawn in any literature (as long as any thing or relationship discovered in the literature is represented in some hierarchy) (Szostak 2013b).

Olson also worries that the logical philosophy that underpins hierarchy privileges reason over emotion and intuition, and assumes away bias. I concur that emotion and intuition are important parts of the process of discovery (Szostak 2002), and have attempted to classify the types of bias that characterize scholarship (Szostak 2004, chapter 5). Classifying works in terms of authorial perspective will, as noted above, provide some insight into potential biases; it may also tell us something about the particular role of emotion and intuition in a work. Allowing free combination will – as in the nurse example above – provide a powerful antidote to bias. And it will be argued in the next section that an emphasis on combinations reduces and may even eliminate the biases that creep into hierarchies themselves.

Though Olson did not describe in detail what her recommended classification would look like, she did appreciate that it would rely heavily on a synthetic approach. She noted that even when a synthetic approach is pursued within contemporary classifications some combinations are privileged over others. It is thus critical that it be possible to freely combine any set of concepts.

Of particular note, Olson appreciates that existing classifications handle paradigmatic relationships best. Yet since paradigmatic relationships are enduring we

rarely need state the obvious. It is syntagmatic relationships (where the connection is not essential, as in embroidery of Christmas ornaments) which we will often wish to express [search for] but these are handled poorly. Boolean searches will yield many hits that do not capture the desired relationship. Again, the solution involves allowing us to freely connect any set of concepts both in classifying a work and in searching.

The sort of classification outlined here, which addresses each of the concerns raised by Olson, not surprisingly serves also the goals identified in the first section. It is universal. It facilitates cross-group exploration and understanding, by relying on combinations of basic concepts that are broadly understood in similar ways across groups. It is amenable to classification by authorial perspective (or group membership) because it allows any concepts to be combined.

Olson is critical not just of classification but of standard practice in constructing thesauri. Hierarchical relationships are captured fairly precisely by the terminology of BT (broader term) and NT (narrower term). But a host of different relationships are lumped together as RT (related term). Our thesauri are thus as guilty of privileging hierarchy as our classification systems. But this need not be: we could aspire to recognizing several key types of relationship. We might in particular designate the basic concepts that combine to generate a more complex concept.

Olson notes that the thesaurus construction standard, ANSI/NISO Z39.19, provides for a limited set of allowed RT relationships: process/agent, process/counteragent, action/property, action/product, action/target, cause/effect, concept or object/property, concept or object/origins, concept or object/measurements, raw material/product, and discipline or field/object or practitioner; and also antonyms (plus a few arcane exceptions) The standard allows these to be explicitly indicated on a local basis. But why not insist that these and others are always designated?[8]

While celebrating diversity is an important (albeit under-studied) desideratum of a classification scheme, it is not the only desideratum. It is thus worth noting that the sort of classification (and thesaurus) urged here has many advantages beyond achieving the goals of the first section and responding to the concerns outlined by Olson. I have argued (2011) that breaking complex concepts into basic concepts is the ambiguity-minimizing strategy in classification.[9] Several other advantages can be briefly noted (see Szostak 2013b):

1. Since most works can be classified as links between phenomena, we are able to achieve very precise classifications with limited and expressive notation.
2. Users are thus better able to find precisely what they want, whether they wish to search in one discipline or across all.
3. By distinguishing different sorts of relationship (especially causation/influence), we enable searches by verb-like terms as well.[10]

---

[8] Olson later suggests that chronological would be one useful addition. Khoo and fin-Cheon 2006 predict that relationships will be of increasing importance in information science. They suggest yet other useful types of relationship that could be indicated including troponymy, intentionality, necessity, conjunction, and disjunction.

[9] Such a classification may also serve as a bridge between other classifications. Yi and Chan (2010) explore the possibility of rendering LCSH interoperable with other systems. They criticize LCSH both for inconsistent application of hierarchy and for unclear semantics and syntax.

[10] Friedman and Smiraglia (2013) find that most concept maps employed in knowledge organization have nouns as nodes and verbs as arcs. But our classifications do not reflect this synergy.

4. While other classification systems provide specific instructions in multiple places for coding by time or place or people, this system has a universal coding for such elements. This renders both classification and searching easier.

5. Note that the use of linked notation serves to place works [but not individual concepts] within multiple hierarchies (and of relations as well as things)

6. It should be possible to translate all search or entry terms employed in other classifications into basic concepts. The system may provide a solution to the fact that online databases employ a bewildering array of classification systems.

7. Note that in addition we create the possibility of (fairly) automatically coding for new works or for existing works that are at present poorly classified.

Lambe (2007) also appreciates that hierarchy is not the only way to classify (he mentions matrices,[11].system maps, and facets), and is often not the best. He appreciates that our goal is to show how things are related. He suggests that users by looking at a taxonomy should gain a sense of how things connect. Taxonomies should also serve as artificial memory aids: helping us to remember things by relating them to others. It deserves to be stressed that a classification that relies on combinations across a very manageable set of schedules is both much easier to master and to understand. Most users approach subject headings within existing classification systems with no understanding of how these are generated or related to each other. The suggested classification is transparent.

Lambe addresses in detail the fact that many concepts appear within multiple hierarchies in existing classifications. He is very critical of this practice, arguing that a hierarchical approach becomes too complicated if concepts appear in many places (see also Soergel 1985, 254-6). It is thus better to capture this sort of situation in other ways than through hierarchy. This is precisely what we have done. And thus our approach serves diversity and also generates a less problematic classification.

It could also be that the approach urged here will close the gap between the fields of classification and information retrieval. Scholars of information retrieval increasingly disdain the 'bag of words' assumption driving many search techniques: that the concepts being searched for occur independently. They appreciate that users search for combinations of concepts (e.g. Mengle and Goharian 2010, Khoo and fin-Cheon 2006). Though search engines generally ignore existing classification systems, they might find a classification which stresses such combinations useful.[12]

Last but not least, Borner (2006) suggests that in the near future scholars might just add 'nuggets' or 'nodes' to the web of knowledge. That is, the present practice of writing stand-alone papers will be replaced by a practice of adding insights to a pre-existing structure. She reviews various efforts over the last century to develop links

---

[11] I developed a five-dimensional typology of theory types in Szostak (2004). This is employed in the Basic Concepts Classification and the Integrative Levels Classification in order to allow us to capture the type of theory employed in a work. See Gnoli and Szostak (2008).

[12] Birger Hjørland and I have often disagreed in the past (see Fox 2012, Szostak 2013c). But in speculating on the role of classification "after Google" he says much that is consonant with the approach recommended here: that information scientists should work on an overall structure that somehow connects domain analyses, that the key is the semantic relations between concepts (though he at times stresses hierarchy), and that documents should be classified not in terms of simple aboutness but rather what a reader would find useful/novel in them.

between related bits of information (such as citation indices). New technology creates an opportunity to finally achieve this goal. But search engines are like inserting a needle in a haystack, and usually do not place search results in context: they "fail to equip scholars with a birds-eye view of the global structure and dynamics of scholarly knowledge and expertise" (186). The sort of classification here can be used both to classify works and ideas (a desiderata noted by Gnoli 2008). It would thus be congenial to the sort of shift foreseen by Borner, such that any author's ideas can readily be related to the ideas of other authors. But the structure's fluidity would mean that classification itself does not privilege certain ideas over others.

Process

The system as outlined above allows the free combination of concepts across any hierarchies (of both things and relationships). How, though, are these hierarchies developed?

Olson has also often used a 'slicing pizza' analogy.[13] We can subdivide classes into subclasses in multiple ways. Different groups may wish to slice their pizza in different ways. This conundrum seems insoluble. Any use of hierarchy in classification must of necessity privilege one way of slicing the pizza.[14]

But it is in fact quite straightforward to address this problem within the sort of classification urged here: most if not all possible approaches to slicing any pizza can in practice be addressed simultaneously within a 'web of relations' approach. As noted above, the web approach significantly lessens the need for hierarchy. In particular, it eliminates the oft-noted practice (e.g. by Mazzocchi et al 2007), common in all major classifications, of abusing hierarchy such that causal arguments (or other sorts of relationship between phenomena) are treated as if they were a proper subset of some phenomenon.

More centrally, when hierarchy is employed in the type of classification recommended here, it is usually subdivision in terms of 'type of' (but occasionally 'parts of') that is called for. And 'type of' is best defined functionally for social phenomena (so that institutions, for example, are classified in terms of their official purpose) and in terms of their essence for natural objects (so that species are organized in terms of genetic inheritance, and chemical compounds in terms of constituent chemicals).

Yet surely this privileges this one slicing strategy? How do we decide that this is the best way to subdivide? While there are other reasons, the one to stress here is that most/all other ways of slicing the pizza can be easily captured through combinations. It has often been noted, for example, that pharmacologists might want to classify drugs in terms of physiological effect, while chemists will want to classify them by chemical composition. The former can easily be rendered as, say (drugs)(reduce)(blood pressure). The latter can only be captured through a 'type of' approach. In other words, the classification wished for by pharmacologists is a classification of relationships, whereas the classification sought by chemists is a classification of subsidiary types of real things.

---

[13] And Olson (2007) appreciates that multiple BT are common in thesauri.

[14] Hjørland (2012) notes there are no good guidelines on how to pursue an interpretivist approach to classification, once one denies there is one best way to subdivide.

The claim here is strong: that there often (always?) is one right way to slice the pizza. We have imagined or at least exaggerated the challenge of slicing because we have abused hierarchy in order to capture relationships. Once we handle relationships as relationships, the slicing conundrum is alleviated and may even disappear. We must always be careful of reaching empirical conclusions on the basis of theoretical arguments alone. The arguments made here must be tested in practice. My own efforts to develop the Basic Concepts Classification suggest that we only rarely confront choices about how to slice (see Szostak 2013b). But this conclusion needs to be verified by others who may bring different slicing preferences to the task. What is clear on theoretical grounds is that we can substantially reduce our slicing choices by treating relationships as relationships.

Mai (2010) argues that information science has long but mistakenly assumed that we were searching for the one best classification, and that general rules and commonalities existed that needed to be identified. As noted above, he argues that bias is inevitable. But it is useful to explore here the precise arguments he makes with respect to what we have termed "slicing the pizza." First, he notes that likeness is not a quality of things but a relationship between them; we can find some similarity between any two things (e.g. plum and lawnmower). But what sorts of similarities exist between a plum and a lawnmower? Perhaps color, perhaps uses to which they can be put, perhaps places they are stored. All of these can be captured through relationships.[15] The only singular class of which they are "types of" is "things." Again, we need to be careful of leaping to an empirical conclusion, but it must seem that many/most/all types of likeness can be handled through a web of relationships approach. And indeed Mai's own words – that likeness is a quality of relationships – suggest that this is so.

Mai then argues, following Hjørland, that a stone in a field has information of different types for different users and thus we cannot hope to classify all of these; no one mapping is the true mapping. But we can clearly use relationships to capture (at least) many of these for they reflect different uses to which the stone can be put: mining, building, skipping, and so on. It is noteworthy that the word "mapping" is used here in the sense of one-to-one mapping when the solution is a map that shows all relationships and thus allows one concept to be mapped to many.

Finally, what about an area where there is intense scholarly controversy, such as in defining types of mental illness? Psychologists disagree about how this is best done (see Cooper 2011). Even here a web of relations approach has much merit. Some psychologists would classify in terms of physiological symptoms and others psychological symptoms. Some would look for common causes, others for common effects. Rather than choosing one way of classifying mental illness, it would be better to employ relationships to capture all.

Concluding Remarks

Though the three issues of purpose, structure, and process were addressed separately, the analyses are complementary: there is one approach to classification that addresses all three. It is thus possible to develop a classification that celebrates and supports diversity. Happily, that approach also has many other positive attributes.

---

[15] Mai (2011) argues that the concept 'cow' can be classified as a food, an animal, and so on. But this is a good example of how we can use relationships to capture its use as a food: animals that can be eaten.

References

Borner, K. 2006. Semantic association networks: Using semantic web technology to improve scholarly knowledge and expertise management. In *Visualizing the semantic web* 2nd ed., ed. V. Geroimenko and C. Chen, Berlin: Springer, pp. 183 -98

Clavier, V. and C. Paganelli 2012. Including authorial stance in the indexing of scientific documents. *Knowledge Organization 39*:4, 292-9.

Cooper, R. 2011. Some classifications will be natural. *Knowledge Organization 38:5,* 398-404.

Feinberg, M. 2011. How information systems communicate as documents: The concept of authorial voice. *Journal of Documentation* 67:6, 1015 – 1037.

Fox, M. 2012. Review of two books by Rick Szostak. *Knowledge Organization 39:*4, 300-3.

Friedman, A., and R.P. Smiraglia. 2013. Nodes and arcs: concept map, semiotics, and knowledge organization. *Journal of Documentation 69*:1, 27 – 48.

Gnoli, C. 2008. Ten long-term research questions in knowledge organization. *Knowledge Organization. 35*: 2-3, 137-149.

Gnoli, C., and R. Szostak. 2008. Classifying by phenomena, theories, and methods: Examples with focused social science theories, In *Culture and identity in knowledge organization,* proceedings of the 10th international ISKO conference, Montréal, 2008, ed. C. Arsenault and J. Tennis, Würzburg: Ergon, pp. 205-211.

Hjørland, B. 2012. Is classification necessary after Google? *Jnl of Documentation* 68:3,299-317.

Lambe, P. 2007. *Organising knowledge: Taxonomies, knowledge and organisational effectiveness.* Chandos.

León Manifesto. 2007. Retrieved April 15, 2013 from www.iskoi.org/ilc/leon.htm

Mai, J. 2011. *The modernity of classification. Journal of Documentation* 67:4, 710-30.

Mai, J. 2010. Classification in a social world: Bias and trust. *Jnl of Documentation* 66:5, 627-42

Mazzocchi, F., Tiberi, M., De Santis, B., & Plini, P. (2007) Relational semantics in thesauri: some remarks at theoretical and practical levels. *Knowledge Organization 34:4,* 197-214.

Olson, H. 2007. How we construct subjects: A feminist analysis. *Library Trends 56:* 2, 509-41.

Mengle, S.S.R., and N. Goharian. 2010. Detecting relationships among categories using text classification. *Jnl of the Amer Society for Information Science & Technology61:*5:1046–1061.

Soergel, D. 1985. *Organizing information: Principles of data base and retrieval systems.* San Diego: Academic Press.

Szostak, R. 2013a. Communicating complex concepts. In *Enhancing Communication and Collaboration in Interdisciplinary Research, ed.* O'Rourke, M., Crowley, S., Eigenbrode, S. D., Wulfhorst, J. D., eds. Thousand Oaks: Sage, pp. 34-55.

Szostak, R. 2013b. *Basic Concepts Classification*. Retrieved April 15, 2013 from: http://www.economics.ualberta.ca/en/FacultyandStaff/~/media/economics/FacultyAndStaff/Szostak/Szostak-Basic-Concept-Classification2.pdf

Szostak, R. 2013c. Speaking truth to power in classification. *Knowledge Organization 40:*1, 76-7.

Szostak, R. 2012. Classifying relationships. *Knowledge Organization* 39:3, 165-78.

Szostak, R. 2011. Complex concepts into basic concepts. *Journal of the American Society for Information Science & Technology* 62:11, 2247-65.

Szostak, R. 2010. Universal and domain-specific classifications from an interdisciplinary perspective. In *Paradigms and Conceptual Systems in Knowledge Organization: Proceedings of the 2010 conference of the International Society for Knowledge Organization* (Rome, February, 2010)*.* eds. C. Gnoli and F. Mazzocchi, Wurzburg: Erlon Verlag, pp. 71-7..

Szostak, R. 2004. *Classifying science: Phenomena, data, theory, method, practice.* Springer.

Szostak, R. 2002. Intuition and interdisciplinarity: A reply to Mackey. *Issues in Integrative Studies* 20, 131-7.

Yi, K., and L.M. Chan. 2010. Revisiting the syntactical and structural analysis of Library of Congress subject headings for the digital environment. *Journal of the American Society for Information Science & Technology* 61:4, 677–687.