

Melanie Feinberg
University of Washington, Seattle, Washington, USA

Beyond retrieval: A proposal to expand the design space of classification

Abstract: In Information Science, the creation of classification schemes has been more commonly described in the mode of scientific discovery, as opposed to artifact design. From the literary warrant of Hulme to the terminological warrant of the Classification Research Group (CRG), to Hjørland's domain analysis, the classificationist seems like one who documents and compiles, and not one who actively shapes design. Outside of Information Science, however, classification is used as an active argument to structure interpretations (in linguistics and philosophy) and as a means of coordinating and imposing order on work practices (in studies of information systems). In this paper, I suggest that classification researchers should investigate a wider variety of design possibilities, in which the purpose of a classification is not assumed to be a retrieval tool in the traditional manner. A consequence of an expanded design space is that standard classification manuals and guides become insufficient support for the design process; a new research area might be the facilitation of problem-setting aspects of classification design. One avenue of possible research involves the description of a design language, following the description of such by Lowgren and Stolterman (2004) and Donald Schon (1983). A design language, which might be operationalized as a set of product qualities, might provide a framework by which designers can better understand, evaluate, and create classifications.

1. Introduction

Within Information Science, the primary purpose of knowledge organization schemes (henceforth generically referred to as "classifications") is often described as information retrieval, or the provision of relevant documents to match a user query (for example, Soergel, 1974; Svenonius, 2003). (The relevant documents might be identified with an indexing term or class number, arranged in advance on a shelf or assembled by request using an online system.) In this paper, I contend that a retrieval emphasis, combined with a prevailing sense that classifications should document some existing state of affairs (as opposed to actively constructing a particular interpretation of the world), unduly limits classification design. Using two examples from outside the literature of knowledge organization, I illustrate additional purposes for classification and suggest that these might be equally appropriate for document classification. Especially as full-text search becomes increasingly effective for basic retrieval tasks, exploring other goals and uses of classification seems not only viable, but necessary.

I propose that future research in classification should address the limitations of current design methods, which focus on the creation of retrieval tools, and consider an expanded design space. While the plurality of possibilities for creating a classification may not allow the specification of precise, standard methods, where one clearly identified step always and inevitably follows another, we may be able to facilitate the design process without rigidly

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

specifying it. One example is to describe a design language, or set of dimensions for the design space, to help classification creators better map their own terrain.

2. Traditions from bibliographic classification: Classification as retrieval tool

The purpose of bibliographic classification is often described as information retrieval. The goal of finding information quickly is a consistent theme: just a few examples include Hulme (who notes that “book classification is a mechanical time-saving operation. . . the discovery of knowledge in books by the shortest route [is] our aim and object”), Foskett (“the aim of an information service is to organise the literature on a systematic basis in order to save the time of research workers”) and Svenonius (“within the discipline of Library and Information Science, CVs [controlled vocabularies] are used primarily to assist in the retrieval of information”) (Hulme, 1911; Foskett, 1974; Svenonius, 2003).

Some might assert that especially the earlier classificationists subordinated the practicalities of retrieval in an attempt to create theoretically sound systems. Bliss, particularly, attempted to provide a theoretical foundation for his library system by using philosophically derived ideas about the order of sciences as the basis for class structure (Bliss, 1929). However, Miksa (1998) notes that despite this overt interest in theoretical justifications from logic and philosophy, library classification has always had a significant pragmatic element. For example, even though the classificationists Richardson and Sayers believed that classifications should follow principles of logic and reflect the “natural order” of the sciences, they were equally practicing librarians who were likely to appeal to common sense as a justification for adjusting the “natural order” so that patrons could more easily locate books on the shelves (Richardson, 1930; Sayers, 1915). Richardson, for example, argues that “the need of adjusting theoretical classification to practical conditions. . . is characteristic of the treatment of all complex concrete things,” and “the object of classification is thus economy and increased efficiency in the use of books.” Miksa (1998) suggests that the retrieval focus of bibliographic classification deepened still further in the 1950s, as the documentation movement, enamored with the needs of specialists and scientists, competed with traditions of librarianship.

If the emphasis on retrieval and classification has been reasonably consistent in the literature, the idea of classification as a reflection of reality (as opposed to an original construction of reality, or as an actively designed product) has been even more influential. In a theoretically derived classification, the goal would be to reflect the true order of all things; however, a more pragmatic classification is also typically seen as a reflection, albeit of a more relative order of reality. The enduring concept of literary warrant directs that a classification should document the world of subjects as exhibited by published works (Hulme, 1911). Even as the idea of general, enumerative classifications that specified classes explicitly lost favor in the mid-twentieth century to “special,” analytic-synthetic (faceted) classifications that presented basic concepts in a particular area, to be combined into complex subjects as necessary, literary warrant remained a key semantic concept in classification development. As Beghtol (1986) notes, literary warrant merely evolved into terminological warrant, and a classification was seen as documenting the way that language is used to represent concepts in a single discipline.

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

While mid-century classificationists such as those involved in the Classification Research Group (CRG) recognized that different subject areas would require different approaches to classification, they did not give up the goal of accurately documenting a sort of true reality. Instead, they attempted to isolate smaller, self-contained “realities” in particular subject areas. For example, a series of books by members of the CRG discusses the need to accommodate different approaches to classification as required by the sciences, the social sciences, and the humanities (Foskett, 1974; Langridge, 1976; Vickery, 1975). In his work on classification in the humanities, Langridge (1976) agrees that “diverse characteristics” of division are required to serve “different interests” in classification. For example, Langridge explains, a rabbit might be characterized as a physical type, as a grassland animal, as a pest, and so on. However, Langridge continues, within the sphere of zoology, rabbits are primarily studied as physical types, which should therefore determine their placement in a zoological classification, even though habitat (and so the rabbit as grassland animal) may also be of interest. While the classification of rabbits may differ according to the discipline under consideration, one should be able within a single discipline to identify a “clearly predominant” characteristic that applies to rabbits; it is the classificationist’s duty to discover that characteristic and document it in the classification. This apparent contradiction—that there are different, equally valid ways to classify and yet that within a particular discipline, a predominant characteristic can be identified—seems tightly connected to the goal of creating a retrieval tool and to the vestigial necessity of creating a linear order for library shelves. If classification’s purpose is to locate information more quickly, then it needs to be commensurate with (what the classificationist sees as) likely user queries. We should put the books on rabbits where most zoologists will expect them.

The domain analytic approach described by Hjørland and Albrechtsen (1995) does not challenge the majority of these assumptions. Hjørland’s primary innovation is to dispute the existence of “clearly predominant” characteristics for classification. For Hjørland, a subject discipline includes a variety of sometimes conflicting approaches to the subject matter. In Hjørland’s (1998) study of the domain of psychology, he describes a domain as constituted by the relationships between phenomena under investigation and epistemological perspectives that shape the research. A domain, for Hjørland, seems to be the union of the approaches used to study it, and so includes various sets of concepts for a single phenomenon, as defined by the included approaches. Psychology, for example, includes both “personality” as defined by behaviorism and “personality” as defined by psychoanalysis. This is certainly a more complex view of a discipline than Langridge describes. However, this just makes for a more complicated structure to document in a classification (which becomes technically possible when we don’t need to think so much about a single place for physical items). Hjørland states that “a psychological classification should represent all the most important approaches and subdisciplines in psychology”; similarly, Abrahamsen’s domain analysis of music indexing concludes that the organization of music should combine the perspectives on musicology that he describes (Hjørland, 1998; Abrahamsen, 2003). While Hjørland(1998) asserts that classification is not neutral and is theory-laden, this seems to be based more on the idea that the material to be classified is theory-laden, than that classificationist is actively designing a certain view in the classification. A domain, for example of psychology, exists; it seems to be the classificationist’s job to find and describe it, not to define or build it.

Of course, there is nothing inherently problematic about the idea of classification as a retrieval tool that achieves its effectiveness through encapsulating the variety of perspectives in a particular domain as faithfully as possible. This may be a perfectly valid and useful strategy to address a particular design situation. I would suggest, however, that because of the messiness that Hjørland describes, the role of the classificationist in shaping the design will be much more active than is typically acknowledged. It is not possible to merely compile a classification without making authorial types of decisions, although one may attempt to keep such intrusions at a minimum. Moreover, as the following two examples from non-bibliographic classification show, other possibilities for design may also exist. Especially as full-text search becomes more effective for basic retrieval tasks, exploring other goals and uses of classification seems imperative.

3. Example 1: Classification as argument

In contrast to Information Science, work from philosophy and cognitive linguistics portrays classification as an argument for a particular interpretation of reality. George Lakoff (1987) contends that the way people categorize is inextricably bound with the way that human beings perceive the world, shaped both by our physical characteristics and by the social structures that we inhabit. Categorization is conditioned by idealized cognitive models (ICMs), which provide a basic template for typical categorization decisions. We have a cognitive model of, for example, what a lie is, and this helps us determine whether or not a particular utterance should be called a lie. But ICMs are incomplete; the world is full of debatable cases. Reasonable people may disagree that the omission of information is a lie. Using an example from Winograd and Flores (1986), if someone says “I’m sorry that I missed the meeting. I had a flat tire,” and yet the flat tire was not actually the cause of missing the meeting, is it a lie? A determination here is influenced by the ICM (and, undoubtedly, one’s role in the communication), but it is equally a *decision* about what the category of lie means in that situation.

Bryant’s (2000) work in philosophy elucidates how even an “expert scientific” classification may be chosen to advance a particular agenda, and yet may also be “true” in that it creates an effective argument for a particular interpretation of reality. One of Bryant’s examples shows how the category “dinosaur” was created by a British scientist, Robert Owen. Previously, dinosaurs had been seen as “fossil lizards,” or simply extremely large reptiles. Owen used a particular set of physical characteristics in the fossil record to argue that, in a certain sense, “dinosaurs” were a separate class, in some ways more related to mammals than to extant reptiles. Bryant notes that previous scientists had seen the same evidence as Owen and yet had not focused on the mammalian characteristics of dinosaurs as opposed to the reptilian ones. Owen did not see more evidence than his predecessors did; he just saw the available evidence differently. Owen, it seems, was particularly motivated to create a new explanation for “dinosaurian” characteristics because he saw the opportunity to use this explanation as a means of discrediting the biological theory of Lamarckism (a form of evolution in which species progressed from simple life forms to the complex perfection of human beings). Bryant contends that Owen was also hoping to discredit a personal enemy, a follower of Lamarck. Owen accomplished both ends by creating a persuasive classification. As the philosopher Broadfield

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

notes, the scientist's "interest is not so much in finding a ground of agreement with other workers on the comparative excellence of various processes, as in proving by a demonstration of the superiority of his own the reasonableness of his dissent" (Broadfield, 1946). Scientific classification progresses via disagreement and argument.

Following Bryant and Broadfield's line of thinking, scientific classification doesn't just document reality; it interprets reality and so creates knowledge. While some Information Science researchers, such as Hjørland, might agree that document classifications interpret reality as well, they are less likely to acknowledge (much less encourage) the possible role of the classificationist in designing this interpretation. The Information Science emphasis on retrieval as the primary purpose of classification can obfuscate the degree to which classifications are designed and not discovered. The characterization of classification as argument emphasizes that the purpose of a classification is not given, but chosen. It also illustrates the extent to which a classification's nature as a document, as being authored (in the sense of being a creative product), can be foregrounded. These properties do not appear restricted to scientific classification, but can be extended to the organization of documents. For example, a historian might classify a collection of archival documents according to personal interpretation of an era, or a gourmet might classify a collection of recipes according to the progression of his or her personal taste. This might be a fascinating way to explore previously unanticipated relationships between items and locate interesting documents (and yet utility as a retrieval tool would undoubtedly be poor, if we take retrieval in the typical restricted sense of matching documents to a query). Indeed, it is my contention that, just as Bryant advocated for perceiving scientific classifications as arguments for interpretations of reality and not as inert reflections of it, so should bibliographic classifications be seen as arguments for interpretations of a subject area and not reflections of how a subject area "really is." When bibliographic classification is created under the currently dominant retrieval paradigm, designers are attempting to make uncontroversial arguments. But a document classification can also make controversial or unique arguments and still be useful. One might find a the perfect recipe for an upcoming soiree in the gourmet's collection, but it might be an unexpected result via an unanticipated path.

Beghtol (2003) compared "naive" classifications to "professional" classifications, claiming that naive classifications are created through scholarly activity for a variety of purposes related to knowledge discovery, while professional classifications are created to facilitate document retrieval. Beghtol recognized similarities of construction between naive and professional classifications, but concluded "that the purposes of naive knowledge discovery classifications differ markedly from the purposes of professional information retrieval classifications," which are "designed to help users find documents." Beghtol's naive classifications are similar to the scientific classifications that I have been discussing, and her professional classifications are document, or bibliographic, classifications. However, where Beghtol sees difference, I see a continuum of similarity. Professional classifications are already to a certain extent like naive classifications—there is already an element of authorship, of creative decision-making, associated with them. In addition, if appropriate for the design situation, there is no reason that the gap between professional and naive classifications could not narrow. Researchers persist, however, in viewing the classification design space as quite limited. In a response to Beghtol, Hjørland and Nicolaisen (2004) object to her use of the word "naive"

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

and claim that, if anything, library classifications are “naive” if they do not accurately reflect scientific and scholarly consensus. Hjørland and Nicolaisen assert that if the field of knowledge organization fails to acknowledge the authority of scientific and scholarly classifications, it will be “totally disrespected and disregarded by the rest of the intellectual world.” I suggest that Hjørland and Nicolaisen see reliance on expert knowledge as necessary primarily because they consider classification in Information Science as building retrieval tools for specialists (where retrieval means matching documents to queries). However, if we allow ourselves to expand the design space for document classification to involve other purposes (such as those that Beghtol describes for “naive” classifications and beyond), then we might readily explore different constructions. Contrary to what Hjørland and Nicolaisen assert, I propose that this does not make knowledge organization unscientific; indeed, as the work of Bryant (2000) emphasizes, the practice of science involves reinterpreting or contradicting received wisdom when it does not fit a particular situation.

4. Example 2: Classification as an element of work practice

Studies of information systems in work contexts show classification as a means of coordinating ongoing projects. While some classifications are used to order collections of documents, the purpose of these classifications is not merely retrieval, particularly in the sense of information seeking commonly employed within Information Science (a user has a query and requires information to satisfy that query). Schmidt and Wagner (2004) describe a variety of classifications in their study of how a Viennese architecture firm coordinates its work activities. As one example, a system for separating CAD drawings into layers and systematically naming each layer enables concurrent, iterative work on building plans. As Schmidt and Wagner relate, one person can work on all the fire escapes in the building while outside consultants simultaneously work on the ventilation shafts. While the layer organization shows regularities, for building projects in general and within particular offices, there is also significant variation according to the dynamic project situation. Creating and evolving the layer organization conventions is “a recurrent concern in the office.” The firm relies on a number of such coordinative artifacts and associated practices, which have a range of ordering functions (besides enabling division of labor and efficient progress, other functions include maintaining consistency of both work processes and work products, enabling version control, and supporting the documentation of regulatory requirements).

As another example, Christiansen (1996) investigates the work of Danish police detectives. Christiansen frames the work of fraud detectives as the compilation of evidence, with the result as a final report on the crime. The fraud detectives that Christiansen studied developed their own database to manage their work. Christiansen describes how the detectives took advantage of the process of organizing information in their database to “systematically walk through and overview the problems of the case.” As work progressed and their sense of the case changed, the detectives would “rearrange the setup so that the menu and key words always reflected their current overall conceptualization of the case.” In addition to helping the detectives organize their ongoing work, use of the database helped the detectives “feel much more confident that everything would be in the final report.”

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

Similarly, in describing their development of document management software for transportation engineers, Trigg, Blomberg, and Suchman (1999) emphasize that a document repository is not merely the software and hardware associated with electronic storage, but a system that involves activities such as preparation and submission of documents, indexing and other metadata creation (and editing of this information as projects evolve), adaptation of the system, and so on. The activities must be coordinated: with other activities, with the people involved, and with other tools in the system (such as the document scanning tools). As described in the work of Schmidt and Wagner, Christiansen, and Trigg, Blomberg, and Suchman, classification is not merely a retrieval tool, nor does its success depend on detailed documentation of the work context; instead, the classification is one element of the ongoing construction of a work context. These studies emphasize both the embeddedness and dynamic aspects of the classifications they discuss. In order to understand these classifications, one needs to understand the people involved with them (classificationists, classifiers, different types of classification users), the processes associated with them, and the other tools that make up the system. When thinking about classification in this way, the design process involves not only what kind of tool should be designed, but the question of how other elements of the system need to adapt. In some sense the work itself enters the design space.

Although recent work in bibliographic classification emphasizes the importance of work context, the focus on classification as retrieval tool seems to limit the perceived design space to exact documentation of the work situation, while the idea of a classification as participant in the work itself seems broader than this. For example, Mai (2004) describes the cognitive work analysis approach to classification design as aiming “to design classification schemes that reflect implicit and explicit invariant structures of the work domain”; according to Mai, “the key notion” of this design approach “is that the classificatory structures in ecological classification schemes reflect the work habits of actors using the systems. . .” If a classification’s purpose is to facilitate ongoing work, it’s not clear that reflecting the specific details of existing practice is the only, or best, way to do so. In the architects’ case, Schmidt and Wagner (2000) assert that the classifications are necessary to impose order on heterogeneous (and otherwise chaotic) local practices; the classifications are constructed out of the work practices of the various collaborators and the more global conventions of the building trades, but don’t merely reflect the practice of any particular group. And yet because specific activities have particular needs, and because work is always changing, the tension between “global ordering systems and the enduring or emerging local concerns and issues” becomes a key element in the design situation. Schmidt and Wagner see a balance between uniformity and flexibility as a key design concern. For example, the architects base project-specific catalogs of design components on the specifications used in reference catalogs, sometimes modifying them for particular needs. Using the standard specifications has several advantages: it takes less time, and the standard specifications are easily understood by outside collaborators (the construction engineers, and so on). Similarly, in their discussion of requirements for infrastructures, Star and Ruhleder (1996) note “the competing requirements of openness and malleability, coupled with structure,” and Suchman (2002) warns that “some appropriation into local circumstances will always be required. The question is not whether that work will need to be done, but with what ease or difficulty.” These assertions seem to indicate that a classification that is too specific and detailed, in a sense too faithful to existing

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

and particular work practices, is not always the best design strategy. For example, it might be more useful, when considering Schmidt and Wagner's architects, to design flexible and extensible core classifications (such as modifying the component catalog to make it easier to adapt in different types of work situations), as opposed to creating an extremely detailed snapshot of current practice.

5. The expanded design space and design research

The two example perspectives from non-bibliographic classification are as different from each other as they are from current conceptions of document classification. A scientific classification presents an original argument; it's not a comprehensive system in the same way as a classification that structures activity between heterogeneous actors. In the design process, very different questions might apply to each. And yet both perspectives might be effectively employed when designing a document classification. These examples show that a classification's purpose cannot be assumed when beginning a design project. Even for the organization of documents, it might be appropriate to design something other than the traditional concept of a retrieval tool.

The classic guides to create bibliographic classifications present their design as a well-defined problem (build a retrieval system) to be solved through a relatively constrained process. Soergel assumes that the decision to build a thesaurus has already been made (Soergel, 1974). Vickery similarly begins with the idea that a classification is needed (Vickery, 1960). Aitchison, Gilchrist, and Bawden briefly discuss only the relative benefits of thesauri and full-text search, and then proceed with instructions for creating a thesaurus (Aitchison, Gilchrist, and Bawden, 2000). These guides can be quite detailed, and include an array of definitions, prescriptive details, and process steps. But they are reticent about how to approach conceptual work, even within the proscribed goal of building a retrieval tool. (For example, Aitchison, Gilchrist, and Bawden may describe what part of speech to use for thesaurus terms, but they don't provide much guidance about how to delineate the overall concept space.) And if we are considering a wider range of design possibilities, without committing to the idea of a retrieval tool, then the guides provide little help.

Part of the difficulty with these guides is a problem-solving orientation. If the problem is not to create a retrieval tool but to figure out what, if anything, to create, then it becomes very difficult to describe the design process in terms of precise methods and process steps. We can't specify the means, because we aren't sure of the ends. This difficulty is also a characteristic exhibited in Rittel and Webber's (1973) idea of "wicked problems": as they state, "the formulation of a wicked problem *is* the problem." Instead of specifying both problem and solution in advance of working on them, Rittel and Webber contend that the problem and its solution should both emerge gradually, in tandem. Similarly, Donald Schon (1983) explains that "when ends are fixed and clear, then the decision to act can present itself as an instrumental problem. But when ends are confused and conflicting, there is as yet no 'problem to solve.'" From this perspective, problem setting (as Schon defines it, "a process in which, interactively, we *name* the things to which we will attend and *frame* the context in which we will attend to them") becomes a key design issue. As an example, Schon describes an architecture professor

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

helping a student. The professor, “Quist,” is able to use domain knowledge of architectural possibilities and forms to have a “reflective conversation with the situation” that the student presents. This conversation opens up a new way of looking at the problem. Quist makes a suggestion, explores the implications of that suggestion, and comes up with new questions for the student to continue working with. Schon describes this as a “successful reframing” of the problem that leads to “the continuance of the conversation.” From this perspective, instead of figuring out how to make a retrieval tool fit a particular situation (solving a predefined problem), the classificationist might contemplate or prototype a solution (“hmmm...what if, instead of creating a single taxonomy for the entire company, we could create a basic template to be customized for different goals?”), see how that illuminates the design problem, and then use this information to decide how to proceed. The search for solutions is also a way to understand the problem better. Schon speaks of *experiments*—a solution is proposed and explored, revealing consequences and implications that provide more insight onto the design situation, in the same way that a scientific experiment provides evidence for or against a hypothesis.

6. Research to facilitate problem setting for classification design

What informs design experiments, and how is a proposed solution to be evaluated? In science, an experiment and the hypothesis it tests are typically grounded in previous work. The research question is presented as a gap in current knowledge, and a predictive hypothesis is made based on existing theories or models, or perhaps on previous findings. In design research, the orientation is similar, except that the idea of previous or related work focuses more narrowly on the individual designer’s skill, experience, and judgment. For example, Vetting Wolf, et al. (2006) criticize the problem-solving focus of “engineering design” for human-computer interaction, noting that many problems benefit instead from “creative design” and an emphasis on problem setting. The practice of creative design that they relate relies heavily on design judgment, or “the ability to assess, appreciate, and make appropriate decisions.” Design judgment pushes the non-linear design process that Vetting Wolf, et al. recount, informs the making of artifacts that they describe as central to progress and communication, and enables the critiques that they feel improve the evolving solution. Certainly, as good design may emerge from a variety of approaches to the situation and may take innumerable forms, there are significant limitations associated with the “engineering” ideal of design as a formalized linear process that proceeds from abstract to concrete, with clear divisions between design and implementation. However, reliance on individual skill, even with the deployment of multidisciplinary teams on design projects, as Vetting Wolf, et al. suggest, is not entirely satisfying as the means to ensure effective problem setting. Not all designers are experts, even those who may have been working in their fields for a long time.

Lowgren and Stolterman (2004) suggest that markers of design expertise include a sense of quality, a language for description, and reflective thinking. The sense of quality pertains to both the product and to the design process. The language for description is a means to express the sense of quality (perhaps through descriptions, perhaps through prototypes). Reflective thinking pulls the sense of quality and language for description into practice. Designers must challenge themselves by reflecting on their own actions (what went wrong? what went right?

how do we move on from this apparent impasse?) and on the work of others (how does that artifact solve a similar problem to the one that I have?). The architect, Quist, that Schon (1983) describes also uses a language of design, in this case specific to architecture. Schon describes Quist's language as including elements such as a building's form, architectural style precedents, and qualities of the site. Quist's fluency in the design language results from his deep experience with architecture. This experience also provides Quist with an extensive "repertoire" of design examples that he can apply or adapt to similar situations.

A possibility for research to facilitate problem setting in classification design is to describe such a design language, one that enables the consistent expression of classification qualities. One can learn a language through immersion alone, but grammars and instruction books make the process easier, faster, and more systematic. Lowgren and Stolterman's idea of product qualities could provide a model for formulating a classification design language. Lowgren and Stolterman claim that few useful attributes of digital artifacts can be measured with scientific evaluation methods. By articulating a more complete set of desirable attributes for digital artifacts (characteristics such as seductiveness, pliability, and transparency), they attempt to provide a framework for designers to better understand, evaluate, and effectively create these artifacts.

To enumerate and describe such qualities for classification design, one might first attempt to analyze and map the design space (acknowledging that it can't be fully mapped), by reviewing both literature about classifications and examining examples of classifications. From this review, a variety of artifacts or descriptions that represent the plurality of design possibilities can be selected, and the qualities that characterize them for good or ill described. Once derived, these qualities can then also be used as a basis for criticism, to show what makes design examples good or bad. Examples that are clearly explicated using the vocabulary of a design language can increase a designer's "repertoire." While using examples is a fairly typical practice in current practitioner guides, such as Rosenfeld and Morville's (2002) primer on information architecture, the examples are typically not explicated or criticized systematically. The way these examples are currently used, one already needs a sense of design judgment to appreciate them.

We might also think of the gradual selection and justification of qualities as a way to build a picture of an evolving design solution, becoming in addition a means to express the design process itself and give it more structure. As described by Floyd, et al. (1989) the work of Peter Naur argues "that the essential thing in developing software. . . is to build a theory on the domain under consideration." Likewise, Rittel and Webber (1973) characterize the design process as "argumentative." A similar motivation underlies the human-computer interaction methods of claims analysis and design rationale, in which decisions and alternatives in the design process are rigorously analyzed and documented (Carroll and Rosson, 1992; Moran and Carroll, 1996). While implementations of design rationale, such as Questions, Options, and Criteria (QOC), are problematic due to their extreme detail and complexity, one can perhaps achieve some of the design rationale goals without the time-consuming documentation required with notations such as QOC (MacLean, et al., 1996; Buckingham Shum, 1996). The idea of using qualities to systematize the design process, as well as the product, is another possible research direction. A means to consistently express product qualities can provide an initial vocabulary to describe what the design should do and why this is a good solution.

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

7. Conclusion

The tradition of Information Science is to perceive classification as a means to retrieve documents that match a user query. This idea has formed the basis of classification design from Hulme to Hjørland, making it appear necessary that a worthwhile classification must accurately document some state of affairs, be it the universe of knowledge, a particular subject area, or the work situation of a company. Examples of classification design outside of Information Science, however, suggest that retrieval tools form just one part of the possible design space. Classifications might be new interpretations of the world, as shown through work in philosophy and linguistics; they might structure and affect work practices, as opposed to passively reflecting them, as shown through work on information systems.

None of this is to say that document classification might not be used for retrieval or should not attempt to replicate a particular knowledge system, if that makes sense for the situation. But a classification might organize documents in a useful or interesting way and yet not be a good retrieval tool. We have no problem thinking of, for example, the organization of artworks in an exhibition as being useful and interesting, perhaps for pedagogical purposes, or perhaps merely for aesthetic pleasure, and yet not suited for retrieval. One might use similar strategies for a Web site about art, or for a collection of art books. Moreover, as automatic search techniques become more sophisticated, the use of classifications for retrieval becomes less necessary and starts to seem anachronistic. Why focus on retrieval when classifications can do more?

However, if one is interested in expanding the classification design space, the current guides and methods provide limited assistance. These manuals present classification work as a defined problem to be solved through the application of precise process steps. Such a problem-solving orientation is insufficient for creative design problems. Instead, creative design stresses the importance of problem setting, of not assuming what you are trying to do, but of actively attempting to figure it out. Here we have a research opportunity for classification: is there a way to facilitate this problem-setting process? In design literature, reliance is typically placed on the skills, knowledge, and judgment of particular designers. Yet this approach seems haphazard. One idea is to develop a design language, perhaps derived from product qualities, as suggested by Lowgren and Stolterman (2004). Articulating such a language can provide a means for designers to communicate, form a base for the development of systematic criticism, and add a measure of structure to the design process.

References

- Abrahamsen, K. T. (2003). Indexing of musical genres. an epistemological perspective. *Knowledge Organization*, 30 (3–4): 144–169.
- Aitchison, J., Gilchrist, A., and Bawden, D. (2000). *Thesaurus construction and use: A practical manual*. 4th ed. London: Europa Publications.
- Beghtol, C. (1986). Semantic validity: concepts of warrant in bibliographic classification systems.” *Library Resources & Technical Services*, 30: 109–125.

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

- Begthol, C. (2003). Classification for information retrieval and classification for knowledge discovery: relationships between “professional” and “naive” classifications. *Knowledge Organization*, 30 (2): 64–73.
- Bliss, H. E. (1929). *The organization of knowledge and the system of the sciences*. New York: Henry Holt.
- Broadfield, A. (1946). *The philosophy of classification*. London: Grafton.
- Bryant, R. (2000). *Discovery and decision: Exploring the metaphysics and epistemology of scientific classification*. Cranbury, NJ: Associated University Presses.
- Buckingham Shum, S. (1996). Analyzing the usability of a design rationale notation. In T. Moran and J. M. Carroll (Eds.), *Design rationale: concepts, techniques, and use*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Carroll, J. M., and Rosson, M. B. (1992). “Getting around the task-artifact cycle: how to make claims and design by scenario. *ACM Transactions on Information Systems*, 10(2): 181–212.
- Christiansen, E. (1996). Tamed by a rose: computers as tools in human activity.” In Bonnie Nardi (Ed.), *Context and consciousness: activity theory and human-computer interaction*. Cambridge, MA: MIT Press.
- Floyd, C., et al. (1989). Out of Scandinavia: alternative approaches to software design and development. *Human Computer Interaction*, 4: 253–350.
- Foskett, D. J. (1974). *Classification and indexing in the social sciences*, 2nd ed. London: Butterworths.
- Hjørland, B. (1998). The classification of psychology: a case study in the classification of a knowledge field. *Knowledge Organization*, 25 (4): 162–201.
- Hjørland, B., and Albrechtsen, H. (1995). Toward a new horizon in information science: domain-analysis. *Journal for the American Society of Information Science*, 46 (6): 400–425.
- Hjørland, B., and Nicolaisen, J. (2004). Scientific and scholarly classifications are not “naive”: a comment to Begthol [sic]. *Knowledge Organization*, 31(1): 55–61.
- Hulme, E. W. (1911). Principles of book classification: chapter II - principles of division in book classification.” *Library Association Record* 13: 389–394.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Langridge, D. W. (1976). *Classification and indexing in the humanities*. London: Butterworths.
- Lowgren, J., and Stolterman, E. (2004). *Thoughtful interaction design: A design perspective on information technology*. Cambridge, MA: MIT Press.
- MacLean, A., Young, R., Bellotti, V., and Moran, T. (1996). Questions, options, and criteria: elements of design space analysis.” In T. Moran and J. M. Carroll (Eds.), *Design rationale: concepts, techniques, and use*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Mai, J-E. (2004). Classification in context: relativity, reality, and representation. *Knowledge Organization*, 31 (1): 39–48.
- Miksa, F. L. (1998). *The DDC, the universe of knowledge, and the post-modern library*. Albany, NY: OCLC.

Feinberg, M. (2007). Beyond retrieval: A proposal to expand the design space of classification. *Proceedings of the North American Symposium on Knowledge Organization*. Vol. 1. Available: <http://dlist.sir.arizona.edu/1892>

- Moran, T., and Carroll, J. M. (1996). Overview of design rationale. In T. Moran and J. M. Carroll (Eds.), *Design rationale: concepts, techniques, and use*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Sayers, W. C. (1915). *Canons of classification*. London: Grafton.
- Schmidt, K., and Wagner, I. (2004). Ordering systems: coordinative practices and artifacts in architectural design and planning. *Computer Supported Cooperative Work*, 11: 349–408.
- Richardson, E. C. (1930). *Classification, theoretical and practical*. 3rd ed. Hamden, CT: Shoe String Press.
- Rittel, H., and Weber, M. (1973). Dilemmas in a general theory of planning. *Policy Science*, 4: 155–169.
- Rosenfeld, L., and Morville, P. (2002). *Information architecture and the World Wide Web*. 2nd ed. Sebastopol, CA: O'Reilly.
- Schon, D. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.
- Soergel, D. (1974). *Indexing languages and thesauri: Construction and maintenance*. Los Angeles: Melville Publishing Company.
- Suchman, L. (2002). Practice-based design of information systems: notes from the hyperdeveloped world. *The Information Society*, 18: 1–6.
- Trigg, R., Blomberg, J., and Suchman, L. (1999). Moving document collections online: the evolution of a shared repository. In S. Bodker, M. Kyng, and K. Schimdt (Eds.), *Proceedings of the Sixth European Conference on Computer-Supported Cooperative Work*, 12-16 September 1999, Copenhagen, Denmark.
- Vickery, B. C. (1960). *Faceted classification*. London: Aslib.
- Vickery, B. C. (1975). *Classification and indexing in the sciences*, 3rd ed. London: Butterworths.)
- Winograd, T. and Flores, F. (1986). *Understanding computers and cognition: A new foundation for design*. Reading, MA: Addison-Wesley.
- Wolf, T. V., et al. (2006). Dispelling design as the “black art” of CHI. In *CHI 2006 Proceedings* (pp. 521–530).