# The role of causality and conceptual coherence in assessments of similarity

**Louise F. Spiteri**
**Louise.Spiteri@dal.ca**
**School of Information Management**
**Dalhousie University**
**Halifax, Nova Scotia. Canada**

## Abstract

Conceptual coherence, which refers to concepts whose contents make sense to the perceiver, has been associated traditionally with the notion of similarity, that is, objects, events, or entities form a concept because they are similar to one another. An examination of traditional similarity-based concept theories suggests that they do not provide an adequate account for conceptual coherence. Library and Information Science needs to explore knowledge-based approaches to concept formation, which suggest that one's knowledge of a concept includes not just a representation of its features but also an explicit representation of the causal mechanisms that people believe link those features to form a coherent whole.

## Introduction

> Concepts are the glue that holds our mental world together …. Concepts tie our past experiences to our present interactions with the world; the concepts themselves are connected to our larger knowledge structures. Our concepts embody much of our knowledge of the world, telling us what things there are and what properties they have (Murphy, 2002, p. 1).

The standard psychological usage of *concept* is that of a mental representation individuated or defined by its contents or features (Laurence & Margolis, 1999). Concepts are tied closely to categories: Categorization involves characterizing something by means of concepts so, for example, my concept of *dog* allows me to pick out a category of entities that I would call dogs (Prinz, 2002).

Conceptual coherence refers to concepts whose contents "seem to hang together, a grouping of objects that makes sense to the perceiver" (Murphy & Medin, 1999, p. 427). Conceptual coherence has been associated traditionally with the notion of similarity, that is, objects, events, or entities form a concept because they are similar to one another. Objects fall into natural clusters of similar kinds (e.g., dogs) that are, at the same time, dissimilar to other kinds (e.g., cats). Concepts are *pattern-recognition devices* that enable us to classify novel entities and to draw inferences about such entities (Smith & Medin, 1981). If I know something about the

properties of the concept *dog*, for example, I am able to make inferences about Spot the Dalmatian, even if I have never encountered a Dalmatian. Thus, "similarity may be the glue that makes a category learnable and useful" (Murphy & Medin, 1999, p. 427). "Concepts give our world stability in that they allow us to treat nonidentical things as equivalent" (Wisniewski, 2002, p. 467).

The notion of similarity, or likeness, underlies most approaches used in Library and Information Science (LIS) in the design of bibliographic classification systems: "Classification is, in its simplest statement, the putting together of like things, or more fully described, it is the arranging of things according to likeness and unlikeness." (Richardson, 1964, p. 1). The reliance upon similarity assumes a shared or common understanding of the attributes or features that give a concept its identity. Does similarity explain, however why a concept was formed or why it makes sense to the perceiver? Will the same concept have the same degree of coherence amongst different people, even within the same domain? Two recent studies that examined how people within the domain of LIS inter-related seemingly similar concepts such as *cataloguing* and *authority control*, for example, showed that although participants agreed that the two terms were similar, they did not agree why they were similar. Some participants said that authority control is a product of cataloguing, while others that cataloguing is a form of authority control (Spiteri, 2004; Spiteri, 2005).

More recently, the influence of the prior *theoretical knowledge* that learners often contribute to their representations of categories has also been a topic of study. For example, people not only know that birds have wings and that they can fly and build nests in trees, but also that birds build nests in trees because they can fly, and fly because they have wings. Many people even believe that morphological features of birds such as wings are ultimately caused by the kind of DNA that birds possess (Rehder, 2003). In comparison, however, with the development of models that account for the effects of similarity and empirical observations, there has been relatively little development of formal models to account for the effects of such prior knowledge.

This paper will examine how conceptual coherence is defined and explored in existing concept theories. It will be argued that traditional similarity-based theories do not provide an adequate

account for conceptual coherence, and that LIS needs to explore newer, knowledge-based approaches to concept formation, which suggest that one's knowledge of many concepts includes not just a representation of a concept's features but also an explicit representation of the causal mechanisms that people believe link those features to form a coherent whole.

**Similarity and cognition**

In the cognitive sciences, similarity is thought to play an essential role in how people acquire and categorize information. Conceptual knowledge involves "the representation of the information required to interact successfully with the environment … the acquisition of conceptual knowledge involves the construction of mental representations that can facilitate that interaction (Hahn & Ramscar, 2001, p. 2). Once knowledge is acquired, similarity plays a

> fundamental role in theories of knowledge and behaviour. It serves as an organizing principle by which individuals classify objects, form concepts, and make generalizations … it is employed to explain errors in memory and pattern recognition, and it is central to the analysis of connotative meaning (Tversky, 1977, p. 327).

Categorization of acquired knowledge proceeds by "comparing new stimuli to previously acquired knowledge representations, and classifying it according to which pre-existing representation it most closely resembles, i.e., according to its similarity to some mental representation" (Hahn & Ramscar, 2001, p. 2). Similarity is crucial to the process of categorization:

> Here is a simple and appealing idea about the way people decide whether an object belongs to a category: The object is a member of the category if it is sufficiently similar to known category members …If you want to know whether an object is a category member, start with a representation of the object and a representation of the potential category. Then determine the similarity of the object representation to the category representation. If this similarity value is high enough, then the object belongs to the category; otherwise, it does not (Rips, 1989, p. 21).

Categorization performs a fundamental function in the process of cognition: "By recognizing similarities between potentially dissimilar entities, the individual is enabled to form theories, or models, of his or her environment that allow him or her to extend to new encounters the generalizations garnered from past experience" (Jacob, 1991, p. 78). Categorization is used also to make predictions about new items that one encounters. If I observe several dogs that have fur, four legs, and a wagging tail, for example, I can conclude that dogs as a category share these

features. Similarity is the foundation of inductive thinking, since categories whose members share similar properties have stronger inductive power than categories whose members are less similar (Heit, 1997; Murphy, 2002). A shared understanding of the nature of a category can facilitate communication. If I say "I have to go home because of my dog" (Murphy, 2002, p. 244), I may not need to explain what I mean by this, as it is generally known that dogs cannot be left alone indefinitely.

**Similarity and bibliographic classification**

The concept of similarity, normally referred to as "likeness" in the LIS literature, is often stated as being an important underlying principle of bibliographic classification. Shera posits four basic assumptions of bibliographic classification: universal order of knowledge; hierarchical (genus-species) division; differentiation; and permanence. More specifically, the principle of differentiation "is derived from the *likeness* or *unlikeness* of the properties or attributes of the component units of the classification" (Shera, 1965, p. 77).

Maltby defines classification as "not only the grouping of things which resemble one another and the separation of those which do not, but the arrangement within each group of its components according to their degree of resemblance" (Maltby, 1975, p. 16). Chan suggests that classification is a process of "deciding on a property or characteristic of interest, distinguishing things or objects that possess that property from those which lack it, and grouping things or objects that have the property of characteristic in common into a class" (Chan, 1994, p. 259).

Richardson argues that likeness "is the universal principle of the order of things … Likeness is so [much part of] the essence of all human thought, that literally there is no smallest part of the human mind which cannot be analyzed into just this operation of distinguishing like and unlike and either holding to or rejecting. Likeness, in particular, is the foundation of that systematic thought carried to its ultimate which we call logic" (Richardson, 1964, p. 6).

Broadfield however, suggests that likeness indicates merely a relationship between things; it is not a characteristic of things. "The quality of things that 'unites' them … is not a likeness nor any kind of relation but a character … the genus, which, being variously differentiated in them,

causes them to be related in various ways" (Broadfield, 1946, p. 2). Likeness thus does not perform a unitary function, but works in concert with genus and species to form a coherent concept. "Resemblance is really only a pointer, indicating the possibility that things might be more profoundly related" (Broadfield, 1946, p. 6).

**Similarity-based theories of concept formation**

*Classical Theory*

The Classical Theory originates largely in classical Greek philosophy, and particularly in the works of Plato and Aristotle, who believed that concepts have a definitional structure whereby they contain necessary and sufficient conditions by which they can be defined. The concept *bachelor*, for example, might be composed of a set of representations such as *is not married*, *is male*, and *is an adult*. Each of these components, or features, specifies a condition that something must meet in order to be a bachelor, and anything that satisfies them all counts as a bachelor, namely, an adult male who is not married; any male who meets all the specified conditions will be a bachelor. Categorization is a process of checking to see if the features that are part of a concept are satisfied by the item being categorized. Presumably, then, entities that are considered similar are members of the same category by virtue of the fact that they share the same properties; thus, different species of the category *mammals*, although they may differ in the sense of being different animals (e.g., dogs, cats, whales), are similar in that they share the properties inherent to mammals. In the Classical Theory, therefore, all members of the same category are equally similar to each other because they possess the same properties; similarity is thus symmetrical, because what is true for one entity in the category is true also for another (Laurence & Margolis, 1999; Rosch, 1999).

*Prototype Theory*

The 1970s saw a shift away from the definitional position of the Classical Theory to the notion that concepts may have a looser structure consisting of features that may not apply equally to all members of the concept. The Prototype Theory argues that all concepts show gradient degrees of membership; for example, a sparrow is a better example of *bird* than is an emu, because a sparrow is associated more readily with the features that one attributes to birds; likewise, fire engine red is a better example of *red* than is red hair (Rosch, 1999; Rosch & Mervis, 1975).

When an object is encountered, its features are mentally represented and compared to prototype representations, which are those items that contain the largest number of typical features (Prinz, 2002). Items can be considered extremely typical, moderately typical, atypical, and borderline concept members—typicality is thus a graded phenomenon. Prototype Theory is tied directly to similarity measures, whereby the judged similarity of any two items is measured by comparing the sets of shared and distinctive features that are associated with them. Prototype Theory is sensitive also to context; for example, dogs or cats might be cited as prototypical pet animals, while elephants and lions prototypical circus animals. Similarly, in Canada, for example, a Canada goose is more likely to be cited as a prototypical bird than would be the case in, say, Germany. The Classical Theory, by contrast, posits that the meaning or definition of a concept should not change according to context (Rosch, 1999).

*Exemplar Theory*

The Classical and Prototype theories both focus upon unitary descriptions that capture the central tendency of any given concept; the difference lies in the acceptance or rejection of a set of necessary and sufficient features to create that description. The Exemplar Theory suggests that people do not have a unitary definition of the concept *dog*, for example, nor is this concept composed of a list of features that is found to varying degrees amongst dogs. Rather, one's concept of *dog* is composed only of the set of dogs that one has actually encountered and remembered; so when I see a Doberman, I compare it to my stored memory of other Dobermans I have encountered. If I encounter an unfamiliar entity, I consult my memory to see which entities it is most similar to; I then add up the perceived similarities to conclude that this new entity (e.g., a Havanese) is also a dog (Smith & Medin, 1999). The Exemplar Theory does not require any defining characteristics; like the Prototype Theory, it deals with prototypes, but rather than rely upon the matching of an entity to a list of salient features, my prototypes are those exemplars with which I am most familiar. So, if I have encountered many Dobermans in my life, this breed would form my prototypical image of the concept *dog*. This means also that my definition of *dog* would be based not on a unitary description that would necessarily apply to a majority of dogs, but to my exemplar of *dog*. In daily life, however, I may have encountered several breeds of dogs, so I might, in fact, have more than one exemplar; for example, a

Doberman may be my exemplar for a large breed of dog, while a Yorkshire Terrier my exemplar for a small breed of dog.

*The insufficiency of similarity to explain concept formation*

In his seminal study of similarity, Goodman concluded that:

> similarity, I submit, is insidious … Similarity, ever ready to solve philosophical problems and overcome obstacles, is a pretender, an impostor, a quack. It has, indeed, its uses, but is more often found where it does not belong, professing powers it does not possess (Goodman, 1972, p. 437).

The crux of Goodman's argument is that saying that two things are similar does not say very much about them, since any two things can be regarded as similar or dissimilar, depending on which features one selects for the purposes of comparison. As Hahn and Ramscar contend, a chair and the pigeon outside the window could share numerous similarities with respect, say, to their closeness to me: "Unless we specify the *respects* in which things are said to be similar, the act of saying that they *are* similar is an empty statement" (Hahn & Ramscar, 2001, p. 3). Over the years, several problems with similarity have been identified, as will be discussed next.

*Circularity*

As has been discussed previously, a coherent concept is one that makes sense to the perceiver: The reason that *bird* is a useful concept is that birds are relatively similar to each other – most birds have wings, lay eggs, and fly, for instance. As Hahn and Chater (1997) suggest, a hypothetical concept *drib* which, grouped with a particular light bulb, Polly the pet parrot, the English Channel, and the ozone layer, would seem to be a highly bizarre concept because the items it groups are not at all similar. Similarity allows us to make generalizations about birds; if we know that Polly is a bird and that she has a beak, it is reasonable to assume that other birds may have beaks. On the other hand, since the constituent members of the *drib* concept are so diverse, it is not reasonable to infer that if one drib has a beak, then so do other dribs.

Items are said to belong to the same concept if they share common properties. The problem is that estimates of similarity may be influenced by people's knowledge that the things being compared are in the same, or different, concepts (Murphy & Medin, 1999). Rather than take the

time to ascertain whether a canary and an emu are actually similar, I assume that they are so because they are both grouped in the *bird* concept. In other words we often see things as being similar because they belong to the same concept, rather than categorizing them *because* of their similarity (Goodman, 1972; Hampton, 1998):

> To say that birds are similar because, among other things, birds generally lay eggs, is the same as saying that birds are similar because, among other things, they are grouped together by the concept *egg-layer*. It seems that objects are similar because they fall under the same concepts. The first point of view suggests that similarity can be used to explain concepts; the second point of view suggests that concepts can be used to explain similarity. This is dangerously circular at best (Hahn & Chater, 1997, p. 44).

*Lack of constraints*

The notion of similarity as a process of matching features or properties has been called into question because it does not provide any constraints on what counts as a feature (Wisniewski, 2002). We say that objects are similar because they have many properties in common. Goodman (1972) questions the number of properties that are required for statements of similarity; a plum and a lawnmower, for instance, can share any number of properties, such as the fact that they weigh less than 100 kilos. All entities can, in fact, have a potentially infinite number of properties in common; likewise, the list of differences could be infinite. Furthermore, there are some attributes that are true of only a small number of the category members – perhaps there are some orange plums, or some lawnmowers run by robots.

> Are two things similar, then, only if they have *all* their properties in common? This will not work either, for of course no two things have all their properties in common. Similarity so interpreted will be an empty and hence useless relation (Goodman, 1972, p. 443).

Any two entities can be deemed similar or dissimilar depending on how many features one uses, and the relevance or salience that one attributes to these features: "Unless one can specify such criteria, then the claim that categorization is based on attribute matching is almost entirely vacuous" (Murphy & Medin, 1999, p. 428). The problem with determining salient features is highlighted by Goodman's "with respect" argument: *X* is similar to *Y* means nothing until it is completed by '*X* is similar to Y with *respect* to property *Z*' (Goodman, 1972). So, if I say that Polly the parrot is similar to an emu, I must specify the properties by which I make this determination.

Goldstone (1994) suggests that a natural reply to the problem of constraints is that all logically-possible properties are not salient or psychologically important. When we make similarity judgments, therefore, we extract and compile features that we consider to be salient for the task at hand; what is considered salient will vary according to the individual and the context. Goodman (1972) uses baggage at an airport check-in as an example of how the salience of features varies by individual. Spectators may notice shape, size, colour, and brand name; the pilot is more concerned with weight, and the passenger with destination and ownership. Which pieces of baggage are more alike than others depends not only upon what properties they share, but upon who makes the comparison, and when.

*Similarity is context dependent*

Goodman's "with respect" argument is overly simplistic, however, because he does not take into account the process of combining properties. People do not usually compare objects only in a single respect, such as size, but along multiple dimensions such as size, colour, shape, and so forth. Given multiple respects, the question is how different factors are combined to give a single similarity judgement. There will be many different similarity values between objects depending on which respects are considered, such as expertise, cultural background, and so forth (Hahn & Chater, 1997; Hampton, 1998; Medin, Goldstone & Gentner, 1993). Different types of similarity can thus be distinguished, depending on the respects in question. The properties that are relevant for a similarity comparison vary widely with age, environment, method of presentation, cerebral hemisphere of processing, level of expertise, and goals; for example, expert and novice physicists evaluate the similarity of physics problems differently, with experts basing similarity judgments more on general principles of physics than on superficial features (Hardiman, Dufresne, & Mestre, 1989; Suzuki, Ohnishi, & Shigemasu, 1992). Similarity may depend also on the context in which a concept is presented (Roth & Shoben, 1983). Tversky (1977) found that when choosing the most similar country to Austria from the set {Sweden, Poland, Hungary}, subjects chose Sweden more often than Hungary. In this case, the dimension *form of government* is important because it highlights a difference between one of the choices (Sweden) and the other two choices. When choosing the most similar country to Austria from the set {Sweden, Norway,

Hungary}, subjects chose Hungary more often than Sweden, because now the feature *Scandinavian* singles out Hungary from the other two candidates.

*Similarity and conceptual coherence*

Similarity-based theories of concept formation do not explain adequately the factors that affect people's perception of why things are alike. One of the underlying assumptions of these models is people's ability to base similarity assessments upon empirical or sensory observations of features; thus, for example, I note that most dogs have a tail that wags, four legs, fur, the ability to bark, and so forth. I can visually distinguish a dog from a cat because, although they may share many of the same attributes (a tail, four legs, and fur), they look different. How reliable are these perceptions of similarity, however? In a much-cited study, Keil (1989) presented pictures of a raccoon that was modified to look like a prototypical skunk to subjects of various ages and found that all subjects over the age of nine insisted that the animal was still a raccoon. The saying "if it walks like a duck and quacks like a duck, it must be a duck" could be countered with the argument that if a person can learn to both walk and quack like a duck, is this person now a duck? The similarity-based theories do not take into account adequately what qualities are essential to *duckness*, but rely only upon a tallying of attributes that constitute a duck. What causes a dog to continue to be a dog, even if a number of typical dog features may not be present? Is my understanding of *dogness* based upon my knowledge of the genetic structure of dogs that makes them unique and different from other animals? How well can I define what constitutes my coherent concept of *dog*? Would I be able to identify all the properties that are unique to dogs, such that *dogness* survives, regardless of any physical alterations? If a dog were genetically altered, would I still regard it as a dog? More importantly, is my understanding of the essence of *dogness* the same as other people's?

The Classical Theory does not account clearly for how we define the essence of an entity, nor, as we have seen, does it account for the fact that some properties may be given different values and weights by different people. The Prototype and Exemplar theories do a better job of acknowledging that context affects what people perceive to be typical examples of a concept, but still do not account for how people define the essence of a concept. Furthermore, the similarity-based theories do not account clearly for the inter-relationships among the properties that

constitute *dogness*, nor how different people combine different features to understand the concept *dog*. Thus, for example, my definition of a dog may rely upon a combination of physical attributes (e.g., the ability to bark) and certain behavioural attributes (e.g., fetching, herding, or retrieving). If I place a higher value on herding behaviour, based perhaps upon my contextual experience and situation, I am likely to associate this attribute more closely with barking, than a person who places a higher value on fetching behaviour, since barking is often an important component of herding. Do some attributes thus have a closer causal relationship with *dogness* than others? Similarity-based theories do not explain sufficiently how our underlying knowledge or understanding of the essence of a concept affects which properties we choose as well as which we combine in causal relationships to affect our understanding of a concept's coherence.

**Knowledge-based models of concept formation**

*Theory-Theory*

Theory-Theory posits that the process of learning about most concepts involves noticing how often properties or features occur and co-occur. Our perception of the salience of features thus depends on how often we encounter them and their correlations, and on our understanding of why these properties co-occur. In Theory-Theory, our formation of concepts is thus influenced by our theories of how features are related; for example, blackness and roundness are both frequently-occurring features of *tyres*, yet roundness seems to be more central to *tyres* since it is so closely linked to the function of tyres (Keil, 2003).

Theory theorists reject similarity-based approaches because of their reliance upon counting attributes whose salience and number cannot be constrained or controlled. Concepts are learned as part of our overall understanding of the world around us. When we learn animal concepts, we integrate this information with our general knowledge of biology, behaviour, and other relevant domains. Concepts are influenced by what we already know, but may serve also to affect our existing knowledge (Murphy, 2002; Rehder, 2003; Rips, 1989); so, for example, recent experiments in the creation of self-replicating robots could cause us to question our current understanding of the biological function of reproduction. Theory-Theory suggests that since concepts should be consistent with what we already know, we use our prior knowledge to decide

whether a new item we encounter belongs in an existing concept, or whether it is necessary to create a new concept.

Theory-Theory believes that we do not rely only on simple observation or feature matching to form concepts: We make inferences based upon our prior knowledge and experience and can add information that we do not observe in the item itself. Returning to Keil's modified raccoon; it continues to be a raccoon even if it is transformed to look like a skunk because the true essence of *skunkhood* lies deeper than superficial qualities. Subjects found that the essence of an entity can include features that are not readily observed, and that even if observable features change, the essence of the entity remains constant. Furthermore, although we may not be able to define what, exactly, is the essence of a raccoon, we presume that it exists (Keil, 1989). This belief in hidden essences is called *psychological essentialism* (Prinz, 2002). Essences are not simply assumed to be defining features, but also the causal reason behind the manifestation of surface features, thus the essence of *tigers* is responsible for all the properties of tigers.

> The causal essentialist bias therefore attributes not only the assumption that many categories have hidden essences, but also the belief that those essences are the reason behind many of the features of a category. The causal essentialist bias does not usually include any sense of how it is that the essence is causally linked to the surface, just the notion that it is (Keil, 2003, p. 673).

*Causal-Mode Theory*

Like Theory-Theory, Causal-Mode Theory accounts for the effects of theoretical knowledge on our understanding of concepts, but places greater emphasis upon causal knowledge, which interrelates or links the features of many concepts that people possess. According to Causal-Model Theory, people's knowledge of many concepts includes not just a representation of a concept's features but also an explicit representation of the causal mechanisms that people believe link those features. People use causal models to determine a new object's category membership (Rehder, 2003).

Causal-Mode Theory is designed to determine the importance, or weight, that individual features have on establishing concept membership. The question of feature weighting is not new, as it has been addressed by Prototype Theory since the 1970s, where weight is often determined by the frequency with which features appear in concept members. Where Causal-Mode Theory differs

is in its focus on how feature weights are determined by people's domain theories; for example, straight bananas are rated as better members of the category *bananas* than straight boomerangs are of the category *boomerangs*, a result people attribute to the default feature *curved* occupying a more theoretically-central position in the conceptual representation of *boomerang* as compared with *banana*. Causal-Mode Theory posits also that particular combinations of features affect people's decision as to what makes for a coherent concept (Rehder, 2003). While both Theory-Theory and Causal-Mode Theory argue in favour of the effect of prior knowledge on perceptions of concept formation, the latter theory places greater emphasis on the role of specific causal knowledge.

**Shortcomings of knowledge-based theories**

A criticism of the knowledge approach is that it does not define clearly what it means by a theory, nor how "attributes, similarity, or context could be derived from theories, either in the abstract or from specific theories. Is this theory of categories as theories a new claim of substance or only a battle cry? What is meant by a theory?" (Rosch, 1999, p. 69). Prinz (2002) argues that with its reliance upon defining essences, the knowledge approach does not provide a sufficient explanation of conceptual structure: If I cannot identify the essence of a raccoon, then how can I have a coherent concept of a raccoon? Would I be able to differentiate, say, between a raccoon and a stuffed raccoon toy? Saying that what makes a raccoon a raccoon is the essence of being a raccoon is circular at best.

Another problem with people's theories is that they could be incorrect or could change over time. People's concept of human evolution could be influenced by whether they accept the theories of creationism versus natural selection; in this case, my concept of human evolution could vary widely from that of another person. Furthermore, the knowledge approach suggests that my understanding of *airplanes*, for example, could be limited or even incomplete unless I have an understanding of physics, aerodynamics, and so forth. Knowledge-based theorists typically allow that people can have rather sketchy theories, but this leads to the inevitable conclusion that they will be unable to form clear concepts. Finally, people's theories may be incorrect, as, for example, the belief that AIDS is the result of divine retribution (Laurence & Margolis, 1999).

The knowledge approach does not address how different people may or may not have the same understanding of a concept. I may believe that *animals* have a soul, while my friend does not. Theory theorists may counter that my friend and I posses the same concept of *animals* because we share similar concept contents; this leads in turn to how we would define similarity without comparing features, but since the knowledge approach does not define concepts in terms of features, then how can we determine whether my friend and I share a common understanding of the *animals* concept? All that the knowledge approach can claim is that my friend and I both have our own *animals* concept.

**The knowledge approach and conceptual coherence**

The criticism that the knowledge approach does not explain adequately how we are able to define the essence, say, of *dog*, is well founded. In my attempt to define a dog, it is probably inevitable that I will list features, attributes, or behaviours that I associate with dogs. Does this mean that I have captured the essence of a dog—perhaps, but that essence could, and is very likely to be, peculiar to me. The similarity approach suggests that my understanding of the concept of *dog* relies upon my ability to list the features that identify uniquely a dog and distinguish it from other concepts. The knowledge approach does not argue against the process of listing features; rather it suggests that one's understanding of what constitutes a dog may include not just a listing of attributes, but also an instinctual understanding of what a dog is, as, for example, "I know a dog when I see one." Furthermore, I may list X number of attributes I associate with a dog, while my friend lists Y number of attributes; our attributes may differ, but does this take away from the fact that both of us are capable of structuring a coherent concept of *dog*? The knowledge approach does not suggest that my friend and I possess the same understanding of *dog*, but states merely that we are capable of forming our own coherent concepts, which reverts to the definition of conceptual coherence, namely, that a concept make sense to the person who forms it.

The knowledge approach does not suggest that there is one, universal essence of doghood. As was discussed above, criticism of the knowledge approach has focused upon its tolerance for (a) individual interpretations of the essence of a concept, (b) the changing nature of one's essence of a concept over time, and (c) incorrect interpretations of the essence of a concept. This criticism

highlights perhaps the most fatal flaw of the similarity approach, namely its assumption that concepts are static, universal, and immutable. It is possible that my interpretation of *doghood* may differ from other people's in details (e.g., I may associate dogs with the chewing of slippers, based on my prior experience, while others may associate dogs with chasing cars), but does this mean that is impossible to achieve a degree of consensus over what constitutes a dog? Rather than insist upon a unitary definition of *doghood*, the knowledge approach acknowledges the existence of *degrees of doghood* that can be agreed upon, especially within a specific domain, and the fact that these areas of consensus may vary across different domains; thus veterinarians' *degrees of doghood* will likely differ from those of, say, animal activists. The knowledge approach posits that we must be willing to accept a degree of uncertainty and some fuzzy boundaries in the design of concepts, but that we can still find enough areas of commonalities to make concepts coherent across a domain.

The fact that one's understanding of a concept may change over time has been noted as a weakness of the knowledge approach, which again points to similarity's assumption that a concept has an existence that is not open to context, environment, or even time itself. The concept of *marriage*, for example, has recently undergone changes in both societal and legal definitions in Canada, such that it no longer necessarily involves the civil union of a man and a woman. Our understandings of concepts change with context, environment, and even personal experience. If anything, the knowledge approach reflects the normal progression of concepts over time that reflects changes in societal and cultural norms.

The third crux of the argument against the knowledge approach rests on its acceptance that one's essence of a concept may be wrong if it is based on erroneous theories of knowledge. Once again, however, the knowledge approach may simply reflect the realities of life. It is possible for one person to accept the tenets of creationism, even if scientific evidence suggests that this theory may be unfounded. Because of its insistence upon personal knowledge, context, and experience, the knowledge approach acknowledges that a person's conceptual coherence may, be fact, be founded upon erroneous information. It is questionable whether similarity can prevent the formation of "wrong" concepts; a person can be provided with all the correct attributes of a concept and still choose to define the concept incorrectly. The question of consensus, however,

may act as a mitigating factor; although not all the members of a domain may agree on the essence of a concept, to the point where some members' interpretation of this concept may be perceived as "wrong," the knowledge approach suggests that it is still possible to establish a baseline level of consensus common to the majority of the members.

**Conclusion: The knowledge approach and bibliographic classification**

The knowledge approach to conceptual coherence parallels recent discussions within LIS about the structure of bibliographic classification systems. Hjørland and Albrechtsen (1999) and Beghtol (2003) argue that classification research must be situated within specific contexts and the domains in which the classification systems are designed to function. Rather than represent a universal truth based upon unitary descriptions of concepts, classification systems represent only particular points of view, whose creators "have to choose to represent one particular view of the knowledge …. [therefore] a classification of a knowledge field … support[s] a given theoretical viewpoint at the expense of other views" (Hjørland & Albrechtsen, 1999, pp. 134-135).

> Any classification is relative in the sense that no classification can be argued to be a representation of the true structure of knowledge … a classification is merely one particular explanation of the relationships in a given field that satisfies a group of people at a certain point in time (Mai, 2004, p. 41).

Classification is based more upon interpretation and judgment than upon logic and its ultimate purpose is to suggest a view of the world that makes sense, or is coherent, to its users (Mai, 2004).

The importance of domain knowledge in the construction of classification systems is reflected in recent LIS studies of domain analysis.

> The domain-analytic paradigm in information science states that the best way to understand information … is to study the knowledge-domains as thought or discourse communities … Knowledge organization, structure, cooperation patterns, language and communication forms, information systems, and relevance criteria are reflections of the objects of the work of these communities and of their role in society (Hjørland & Albrechtsen, 1995, p. 400).

The socio-cognitive view of domain analysis emphasizes how domains structure culturally-produced signs and symbols and how its members mediate their cognitive processes into coherent concepts that reflect shared cultural, historical, and social meanings. Members of a

domain have both individual knowledge structures and biases, as well as shared views: "There is an interplay between domain structures and individual knowledge, an interaction between the individual and the social level" (Hjørland, 2004, p. 409).

The knowledge approach's emphasis upon consensus has parallels also within LIS. Consensus in classification can be traced to Henry Evelyn Bliss (1939), who believed that classification systems should reflect how members of the scientific and educational communities structure knowledge in their respective domains. Bliss believed that scientific and educational consensus tends to become permanent and is therefore a sufficiently stable basis for a general bibliographic classification system (Beghtol, 1995). Hjørland and Albrechtsen, however, argue that:

> today, it is regarded as somewhat naive to think that consensus guarantees truth … [but] this does not automatically reject consensus building as a method. In fact, an important characteristic of a subject area might actually be its degree of stability, degree of consensus among the researchers at a given time (Hjørland & Albrechtsen, 1995, p. 402).

There is an opportunity for consensus to play a potentially important role in the major bibliographic subject access and classification systems. The Dewey Decimal Classification (DDC) Editorial Policy Committee, for example, consults regularly with subject experts and members of the LIS community "on matters relating to changes, innovation and the general development of the DDC" (OCLC, 2005). The Library of Congress established the Subject Authority Cooperative Program to provide a means for libraries to submit subject headings and classification numbers to the Library of Congress via the Program for Cooperative Cataloging (Library of Congress, 2005).

The dependence of many LIS bibliographic classification systems upon similarity-based measures of conceptual coherence may result in systems that impose a unitary definition of coherence on any given concept. As has been shown, similarity-based measures do not account adequately for the effect of factors such as context, environment, time, culture, society, knowledge, and expertise upon the definition and coherence of concepts, which is in stark contrast to the knowledge approach, which assumes that these factors do, in fact, affect the development of coherent concepts. The knowledge approach is based also on the assumption that although a unitary definition of a concept is probably not possible, or even desirable, it may still be possible to find enough consensus within knowledge domains to create concepts that are

coherent to the majority of people, with the understanding that this consensus is prone to change and may be reflective of only a given context in a given time.

Given the demonstrated shortcomings of the similarity approach, there is a need within LIS to examine the impact of knowledge and causality upon people's construction of concepts and to see whether it is possible to achieve a *consensus of coherence* for these concepts within given domains. The Causal-Model theory is the most developed working model of the knowledge approach in its formal account of how causal knowledge influences the importance of features and specific configurations of features in judgments of concept membership. This model's ability to provide a precise, quantitative account of both the differences in feature weights and the importance of feature configurations induced by people's knowledge thus makes it an attractive candidate for integrating the knowledge approach into the construction of bibliographic classification systems.

**References**


Beghtol, C. (1995). Domain analysis, literary warrant, and consensus: The case of fiction studies. *Journal of the American Society for Information Science*, 46(1), 30-44.

Beghtol,C. (2003). Classification for information retrieval and classification for knowledge discovery: Relationships between "professional" and "naïve" classifications. *Knowledge Organization*, 30(2), 64-73.

Bliss, H. E. (1939). *The organization of knowledge and the subject-approach to books*. New York: H. W. Wilson.

Broadfield, A. (1946). *The philosophy of classification*. London: Grafton & Co.

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

Goldstone. R. (1994).The role of similarity in categorization: Providing a groundwork. *Cognition*, 52, 125-157.

Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects* (pp. 437-447). Indianapolis: Bobbs-Merrill.

Hahn, U., & Chater, N. (1997). Concepts and similarity. In K. Lambert & D. Shanks (Eds.). *Knowledge, concepts and categories* (pp. 43-92). Cambridge, MA: MIT Press.

Hahn, U., & Ramscar, M. (2001). Introduction: Similarity and categorization. In U. Hahn & M. Ramscar (Eds.), *Similarity and categorization* (pp. 1-11). New York: Oxford University Press.

Hampton, J. A. (1998). Similarity-based categorization and fuzziness of natural categories. *Cognition*, 65, 137-165.

Hardiman, P. T., Dufresne, R., & Mestre, J. P. (1989). The relation between problem categorization and problem solving among experts and novices. *Memory & Cognition*, 17, 627-638.

Heit, E. (1997). Features of similarity and category-based induction. In Ramscar, M.; Hahn, U.; Cambouroplos, E.; & Pain, H. (Eds.), *Proceedings of the interdisciplinary workshop on categorization and similarity, University of Edinburgh* (pp. 115-121). Retrieved January 11, 2007, from http://faculty.ucmerced.edu/eheit/simcat.pdf

Hjørland, B. (2004). Domain analysis: A socio-cognitive orientation for information science. *Bulletin of the American Society for Information Science and Technology*, 30(3), 17-21.

Hjørland, B., & Albrechtsen, H. (1995). Toward a new horizon in information science: Domain-Analysis. *Journal of the American Society for Information Science*, 46(6), 400-425.

Hjørland, B., & Albrechtsen, H. (1999). An analysis of some trends in classification research. *Knowledge Organization*, 27(1), 131-139.

Jacob, E. K. (1991). Classification and categorization: Drawing the line. In B. H. Kwasnik & R. Fidel (Eds.), *Advances in classification research, volume 2. Proceedings of the 2nd ASIS SIG/CR classification research workshop* (pp. 67-83). Medford, NJ: Learned Information, Inc.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Keil, F. C. (2003). Categorisation, causation, and the limits of understanding. *Language and Cognitive Processes*, 18(5/6), 663-692.

Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. *3*-81). Cambridge, MA: MIT Press.

Library of Congress. (2005). *About the SACO program: Program for cooperative cataloging*. Retrieved January 11, 2007, from http://www.loc.gov/catdir/pcc/saco/sacopara.html.

Mai, J. E. (2004). Classification in context: Relativity, reality, and representation. *Knowledge Organization*, 31(1), 39-48.

Maltby, A. (1975). *Sayers' manual of classification for librarians* (5th ed.). London: André

Deutsch.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100, 254-278.

Murphy, G. L. (2002). *The big book of concepts.* Cambridge, MA: MIT Press.

Murphy, G. L., & Medin, D. L. (1999). The role of theories in conceptual coherence. In E. Margolis & S. Laurence (Eds.), *Concepts: Core reading* (pp. 425-458). Cambridge, MA: MIT Press.

OCLC. (2007). *About DDC*. Retrieved January 11, 2007, from http://www.oclc.org/dewey/about/epc/.

Prinz, J. J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.

Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1141-1159.

Richardson, E. C. (1964). *Classification* (3rd ed.). Hamden, CT: Shoe String Press.

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosnaidu & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). Cambridge: Cambridge University Press.

Rosch, E. (1999). Reclaiming concepts. *Journal of Consciousness Studies*, 6: 61-77.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.

Roth, E. M., & Shoben. E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15, 346-378.

Shera, J. H. (1965). *Libraries and the organization of knowledge*. Melbourne: F. W. Cheshire.

Smith, E. E., & Medin. D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Smith, E. E., & Medin. D. L. (1999). The exemplar view. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 207-221). Cambridge, MA: MIT Press.

Spiteri, L. F. (2004). Word association testing and thesaurus construction. *Libres*, 14(2). Retrieved January 11, 2007, from http://libres.curtin.edu.au/libres14n2/index.htm

Spiteri, L. F. (2005). The use of word association in the construction of information retrieval thesauri: A pilot study. *Cataloging & Classification Quarterly*, 40(1), 55-78.

Suzuki, H., Ohnishi, H., & Shigemasu, K. (1992). Goal-directed processes in similarity judgment. *Proceedings of the fourteenth annual conference of the Cognitive Science Society* (pp. 343-348). Hillsdale, NJ: Erlbaum.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352.

Wisniewski, E. J. (2002). Concepts and categorization. In H. Pashler & D. Medin (eds.), *Stevens' handbook of experimental psychology. Third edition. Volume 2: Memory and cognitive processes* (pp. 467-530). New York: John Wiley & Sons.