2000

The role of classification in knowledge represantation and discovery

Barbara H. Kwasnik, *Syracuse University*

The Role of Classification in Knowledge Representation and Discovery

BARBARAH. KWASNIK

ABSTRACT

The link between classification and knowledge is explored. Classification schemes have properties that enable the representation of entities and relationships in structures that reflect knowledge of the domain being classified. The strengths and limitations of four classificatory approaches are described in terms of their ability to reflect, discover, and create new knowledge. These approaches are hierarchies, trees, paradigms, and faceted analysis. Examples are provided of the way in which knowledge and the classification process affect each other.

INTRODUCTION

Developments in our ability to store and retrieve large amounts of' information have stimulated an interest in new ways to exploit this information for advancing human knowledge. This article describes the relationship between knowledge representation (as manifested in classifications) and the processes of knowledge discovery and creation. How does the classification process enable or constrain knowing something or discovering new knowledge about something? In what ways might we develop classifications that will enhance our ability to discover meaningful information in our data stores?

The first part of the article describes several representative classificatory structures-hierarchies, trees, paradigms, and faceted analysis with the aim of identifying how these structures serve as knowledge representations and in what ways they can be used for knowledge discovery and creation. The second part of the discussion includes examples from existing classification schemes and discusses how the schemes reflect or fail to reflect knowledge.

KNOWLEDGE, THEORY, AND CLASSIFICATION

Scholars in many fields, from philosophy to cybernetics, have long discussed the concept of knowledge and the problems of representing knowledge in information systems. The distinction is drawn between merely observing, perceiving, or even describing things and truly knowing them. To know implies a process of integration of facts about objects and the context in which the objects and processes exist. Even in colloquial usage, knowledge about someone or something is always expressed in terms of deep relationships and meanings as well as its place in time and space. To know cars means not only understanding car mechanics but also knowledge of the interplay of the mechanical processes and perhaps even factors such as aesthetics, economics, and psychology.

The process of knowledge discovery and creation in science has traditionally followed the path of systematic exploration, observation, description, analysis, and synthesis and testing of phenomena and facts, all conducted within the communication framework of a particular research community with its accepted methodology and set of techniques. We know the process is not entirely rational but often is sparked and then fueled by insight, hunches, and leaps of faith (Bronowski, 1978). Moreover, research is always conducted within a particular political and cultural reality (Olson, 1998). Each researcher and, on a larger scale, each research community at various points must gather up the disparate pieces and in some way communicate what is known, expressing it in such a way as to be useful for further discovery and understanding. A variety of formats exist for the expression of knowledge e.g., theories, models, formulas, descriptive reportage of many sorts, and polemical essays.

Of these formats, science particularly values theories and models because they are a “symbolic dimension of experience as opposed to the apprehension of brute fact” (Kaplan, 1963, p. 294**)** and can therefore be symbolically extended to cover new experiences. A theory thus explains a particular fact by abstracting the relationship of that fact to other facts. Grand, or covering, theories explain facts in an especially eloquent way and in a very wide (some would say, universal) set of situations. Thus, Darwinian, Marxist, or Freudian theories, for example, attempt to explain processes and behaviors in many contexts, but they do so at a high level of abstraction. There are relatively few grand theories, however, and we rely on the explanatory and descriptive usefulness of more “local” theories-theories that explain a more limited domain but with greater specificity.

CLASSIFICATION AS KNOWLEDGE REPRESENTATION

How are theories built? How does knowledge accumulate and then get shaped into a powerful representation? There are, of course, many processes involved, but often one of them is the process of classification. Classification is the meaningful clustering of experience. The process of classification can be used in a formative way and is thus useful during the preliminary stages of inquiry as a heuristic tool in discovery, analysis, and theorizing (Davies, 1989). Once concepts gel and the relationships among concepts become understood, a classification can be used as a rich representation of what is known and is thus useful in communication and in generating a fresh cycle of exploration, comparison, and theorizing. Kaplan (1963) states that “theory is not the aggregate of the new laws but their connectedness, as a bridge consists of girders only in that the girders are joined together in a particular way” (p. 297). A good classification functions in much the same way that a theory does, connecting concepts in a useful structure. If successful, it is, like a theory, descriptive, explanatory, heuristic, fruitful, and perhaps also elegant, parsimonious, and robust (Kwasnik, l992b).

There are many approaches to the process of classification and to the construction of the foundation of classification schemes. Each kind of classification process has different goals, and each type of classification scheme has different structural properties as well as different strengths and weaknesses in terms of knowledge representation and knowledge discovery. The following is a representative sample of some common approaches and structures.

HIERARCHIES

We have inherited our understanding of hierarchical classifications from Aristotle (Ackrill, 1963), who posited that all nature comprised a unified whole. The whole could be subdivided, like a chicken leg at the joint, into “natural” classes, and each class further into subclasses, and so on-this process following an orderly and systematic set of rules of association and distinction. How do we know what a natural dividing place is, and how do we arrive at the rules for division and subdivision? According to Aristotle, only exhaustive observation can reveal each entity’s true (essential) attributes, and only philosophy can guide us in determining the necessary and sufficient attributes for membership in any given class. In fact, according to Aristotle’s philosophy, it is only when an entity is properly classed, and its essential properties identified, that we can say we truly know it. This is the aim of science, he claims-i.e., to unambiguously classify all phenomena by their essential (true) qualities.

While Aristotle’s legacy is alive in spirit in modern applications of classification, most practitioners recognize that a pure and complete hierarchy is essentially possible only in the ideal. Nevertheless, in knowledge domains that have theoretical foundations (such as germ theory in medicine and the theory of evolution in biology), hierarchies are the preferred structures for knowledge representation (see, for example, the excerpt from the Medical Subject Headings [MeSH] in Figure 1.

- Eye Diseases
  - Conjunctival Diseases
    - Conjunctival Neoplasm
    - Conjunctivitis
      - Conjunctivitis, Allergic
      - Conjunctivitis, Bacterial
        - Ophthalmia Neonatorum
        - Trachoma
      - Conjunctivitis, Viral
      - Keratoconjunctivitis
      - Reiter’s Disease
  - Corneal Diseases
    - Etc.

Figure 1. Hierarchy: Excerpt from MeSH (Medical Subject Headings)

Based on the MeSH excerpt in Figure 1, note that hierarchies have strict structural requirements:

Inclusiveness: The top class (in this case, **EYEDISEASES)** is the most inclusive class and describes the domain of the classification. The top class includes all its subclasses and sub-subclasses. Put another way, all the classes in the example are included in the top class: **EYEDISEASES.**

Species/differentia: A true hierarchy has only one type of relationship between its super- and subclasses and this is the generic relationship, also known as *species/differentia*, or more colloquially as the *is-a* relationship. In a generic relationship, **ALLERGIC CONJUNCTIVITIS** ***is*** *a* kind of **CONJUNCTIVITIS,** which in turn ***is*** *a* kind of **CONJUNCTIVAL DISEASE,** which in turn ***is*** *a* kind of **EYEDISEASE.**

Inheritance: This requirement of strict class inclusion ensures that everything that is true for entities in any given class is also true for entities in its subclasses and sub-subclasses. Thus whatever is true of **EYE DISEASES** (as a whole) is also true of **CONJUNCTIVAL** Whatever is true of **CONJUNCTIVAL DISEASES** (as a whole) is also true of **CONJUNCTIVITIS,** and so on. This property is called *inheritance*, that is, attributes are inherited by a subclass from its superclass.

Transitivity: Since attributes are inherited, all sub-subclasses are members of not only their immediate superclass but of every superclass above that one. Thus if **BACTERIAL CONJUNCTIVITIS** is a kind Of **CONJUNCTIVITIS,** and **CONJUNCTIVITIS** is a kind Of **CONJUNCTIVAL DISEASE,** then, by the rules of transitivity, **BACTERIAL CONJUNCTIVITIS** is also a kind of **CONJUNCTIVAL DISEASE,** and so on. This property is called transitivity.

Systematic and predictable rules for association and distinction: The rules for grouping entities in a class (i.e., creating a species) are determined beforehand, as are the rules for creating distinct subclasses (differentia). Thus all entities in a given class are like each other in some predictable (and predetermined) way, and these entities differ from entities in sibling classes in some predictable (and predetermined) way. In the example above, **CONJUNCTIVAL** and **CORNEAL DISEASES** are alike in that they are both kinds of EYE**DISEASES.** They are differentiated from each other along some predictable and systematic criterion of distinction (in this case "part of the eye affected")

Mutual exclusivity: A given entity can belong to only one class. This property is called mutual exclusivity.

Necessary and sufficient criteria: In a pure hierarchical classification, membership in a given class is determined by rules of inclusion known as necessary and sufficient criteria. To belong to the class, an entity must have the prescribed (necessary) attributes; if it has the necessary attributes, this then constitutes sufficient warrant, and the entity must belong to the class.

Because of these formal properties, hierarchical classification schemes continue to have great appeal in knowledge representation and discovery for several reasons:

Complete and comprehensive information: A hierarchical classification is usually a fairly comprehensive classification since all rules for aggregation and distinction must be made a priori. This means that, before the structure is established, the designer must know a great deal about the extent of the entities, their attributes, and the important criteria along which they are similar and different.

Inheritance and economy of notation*:* The formalism of a hierarchy allows an economical representation of many complex attributes. Each attribute does not have to be repeated at each level but rather is inherited as part of the scheme. Much information can be "carried" by the hierarchical structure.

Inference*:* For this reason, a hierarchy allows reasoning from incomplete evidence. If it can be established, for instance, that a patient has the symptoms of conjunctivitis (as defined by the necessary and sufficient criteria by which a set of symptoms is given this label), then it is possible to know also that, as a kind of eye disease, conjunctivitis will share properties with other eye diseases. This is especially useful if the shared criteria are not obvious or easily observable. For example, if, by observation and comparison with other animals, you assess that an animal is a kind of *cat,* which is a kind of *mammal,* you can infer and predict that, if it is a female, it will reproduce by bearing live young and breast feeding its babies, even though these cat-like characteristics may not be immediately evident.

Real definitions: Hierarchical classification enables *real definitions,* which are considered by many to be superior to other types of definitions because they provide a way of expressing how an entity is *like* something, and also how it is *different* in some important way. For instance, consider the definition: “A bachelor is an unmarried man.” A bachelor is a man; therefore he shares all the characteristics of men. Men can be married or unmarried. A bachelor is of the “unmarried type of man. The strength of this definition, as a definition, lies in its ability to succinctly describe a complex of attributes of affinity and an important aspect of distinction. Two alternative definitional strategies to real definitions are to list attributes one by one or to point to exemplars (“See that guy? He’s a bachelor.” “See that other guy? He’s a bachelor too”). A real definition is often the more efficient way of describing the nature of the entity and the boundaries of where, by definition, that entity ends.

High-level view and holistic perspective: If the criteria by which the classificatory structure is built are theoretical in the sense that they reveal fundamental and meaningful distinctions, then the classification scheme as a whole provides a visualization of the phenomena it is representing. Such a bird’s eye perspective enables recognition of over- all patterns and anomalies, interesting or problematic relationships, and so on. A holistic high-level view is often a trigger for knowledge generation, allowing the researcher to step away from the individual instances to see them as they fit into a larger context.

Not every knowledge domain lends itself to being represented by a hierarchy, however. While hierarchies are desirable for their economy of notation, the richness of description, and the incorporation of knowledge about relationships, they are also problematic for a number of reasons:

Multiple hierarchies: At the top of the list is the fact that, from our modern (non-Aristotelian) perspective, we no longer view the world as having only *one* reality-i.e., one way of being parsed neatly at the joints. Most phenomena are understood to have several, perhaps over- lapping, but separate sets of attributes and relationships, depending on the context and goal of the representation. For instance, dogs are mammals and knowing they are mammals helps us understand their physiological selves. But dogs are also pets and as such belong to the domain of domesticated animals and human companions. Knowing this aids in understanding the social aspects of dog behavior in a particular context, as well as the larger social phenomenon of pet ownership. This suggests that we must have separate classifications for “dogs as animals” and “dogs as pets” with perhaps some cross-links to show the connections in a tangled, or multi hierarchical, structure. In any event, no one classification is able to capture all aspects of a particular domain.

Multiple and diverse criteria: There seem to be some practical limits to how much information a hierarchy can bear in its structure before it becomes too complex. Consider the placement of *lions* in a classification of animals. Traditional zoological taxonomy, based on morphological attributes, places *Lions* in with other *felines.* But consider the distinction between lions in the wild versus lions in zoos. Are they the same entity? A hierarchy is not well designed to accommodate distinctions made along two very different sets of criteria. While it is possible in theory to further subdivide each animal in taxonomy of animals by whether it is in the wild or in captivity, such a representation becomes very cumbersome and repetitive. If a hierarchy is weighted down by too many perspectives and disparate rules for grouping and differentiation, it loses some of its power as a clear representation. One of the difficulties with traditional taxonomies of the living world, in fact, is its inability to accommodate the notion of “habitat.” The representation of knowledge about living entities in ecological systems and over time is difficult in a hierarchy that requires conformity to the principle of mutual exclusivity. For example, in classifying *dinosaurs,* one must decide whether it is more useful to cluster a particular dinosaur under the domain of *prehistoric creatures* (thus using “age” as the defining factor) or to separate dinosaurs and classify each particular kind under the domain of *mammals, birds, reptiles,* and *so* on (thus focusing on their attributes as specific types of animals rather than on when they lived). To do both simultaneously is representationally difficult.

Lack of complete and comprehensive knowledge: Since hierarchies attempt to be comprehensive and to show the relationship of all entities to each other in an overall structure, they require relatively complete knowledge of the domain in advance. In emerging fields, where the extent of the domain is not yet charted, where the relationships are not yet fully understood or defined, or where there is no theoretical framework on which to build the structure, a hierarchy is both difficult and inappropriate to build.

It is not just a question of comprehensiveness. If a knowledge domain rushes into a hierarchical representation without adequate grounding or warrant, the result can be a representation that is misleading or skewed. Such representations can also lead to premature closure in terms of knowledge creation because a hierarchy implies clear boundaries and a complete set of criteria, while this may not in fact be the case. The sure sign of a "premature" hierarchical structure is the need for a "miscellaneous" or "other" category into which the classifier places all those entities that do not fit into the logic of the classification system as specified.

Differences of scale: In order to maintain the principles of transitivity and inheritance, all entities in a hierarchy must be at the same conceptual level of granularity. For example, in classifying the entity *beach,* it is possible to look at a beach from the global perspective and see *"an area* ***of*** *demarcation between land and sea,"* or from the perspective of a human walking on it as *"sand, shells, seaweed, etc.* "or through a micro- scope as *"crystalline structures"* -same beach, different level of definition. Such differences in scale are not easy to accommodate in one classification. If combined into one structure, and especially if combined haphazardly, they weaken the integrity of the knowledge representation. This is because it is not clear at any given point in the classification which criteria of association and distinction are being invoked: beach as *land-mass type,* beach as *habitat,* or beach as *physical material.*

Lack of transitivity: A hierarchy requires that attributes are passed on down the structure intact. So, if *A* is a subclass of B, and *B* is a subclass of C, then A is also a subclass of C. This neatness does not always translate into the way we humans perceive the phenomena around us. For example, we might all agree that *chairs* are a kind of *furniture.* Further, we might agree that rocking *chairs,* and *easy chairs,* possibly *stools,* and perhaps even *tree stumps* are a kind of chair-depending on the context. But, while most people would agree that a *stool* is also, therefore, a kind of *furniture* (thus conforming to the principle of transitivity), not all people would extend the inheritance and agree that a *tree stump* is a kind of *furniture.* In other words, somewhere in the chain of representation the rules change and not all the attributes of *furniture* get invoked in determining the nature of a *tree stump.* This situation leads to a knowledge representation that subtly shifts. As a consequence, it is not possible to use such a representation as a reliable source of inference.

Rules for class inclusion are too strict: Entities do not always conform to the *necessary-and-sufficient* criterion. In a pure hierarchy, entities must belong unambiguously to a class. If they possess all the necessary attributes, they are in; if they lack any of the attributes, they are precluded from membership. In a hierarchy, each member of a class is therefore as good a representative of its class as any other. Unfortunately, human beings do not perceive things quite so neatly. Entities can belong to a class more or less. The criteria for inclusion might fit one entity better than they do another. One entity might be a better representative of a class than another. For instance, most people think a *robin* is closer to the prototype of a *bird* than is a *penguin.* Put another way, *penguins* are not unambiguously a member of the *bird* class, even though they may in fact possess all the necessary and sufficient at- tributes of bird-ness but not to the same perceived degree as a *robin* does. Furthermore, entities in a class may share some attributes in common with each other, but not all might share the *same* attributes. Thus, in my family, there may be a distinctive nose, distinctive eye- brows, and a distinctive smile, but not all members must have all these attributes to be perceived as showing a family resemblance. Finally, an entity may belong to one class under one set of circumstances, and to another class under another set of circumstances, or to both simultaneously. One can be both a *parent* and a *student* or sometimes a *parent* and sometimes a *student.* It is possible to be sometimes a better exemplar of a *student* (closer to the prototype), while at other times less prototypical. This fuzziness requires a different method of representation-some mechanism for indicating relative weight and presence of attributes and relative closeness or distance from some best example prototype. With permeable membranes and dynamic membership in classes, it is difficult to maintain the principles of transitivity and inheritance.

In summary, hierarchies are excellent representations for knowledge in mature domains in which the nature of the entities, and the nature of meaningful relationships, is known. Hierarchies are useful for entities that are well defined and have clear class boundaries. In general, some theory or model is necessary to guide the identification of entities, the rules of association and distinction, and the order in which these rules are invoked.

Trees: Another type of classificatory structure used to represent entities and their relationships is a tree. A tree divides and subdivides its classes based on specific rules for distinction just as in a hierarchy but does not assume the rules of inheritance. Thus, in a tree, the entities have systematic relationships but not the generic (is-a) relationship. There are many types of relationships that can he represented by a tree (see, for example, Figure 2).

GENERALS
COLONELS
CAPTAINS
LIEUTENANTS
SERGEANTS
PRIVATES

Figure 2. Tree: Chain of Command in the Army.

In this tree, the entities are the names of Army ranks. The relationships among the ranks can be described as "chain of command" or "who re- ports to whom." That is, a GENERAL, *commands COLONELS, COLONELS report to GENERALS, COLONELS command CAPTAINS* and *so* on. GENERALS *command* PRIVATES as well, although not directly, but a PRIVATE is not a kind of SERGEANT, and a SERGEANT is not a kind of LIEUTENANT, so the principle of division by *species/differentia* does not apply. Conversely, SER- GEANTS do not inherit the attributes of LIEUTENANTS. In terms of knowledge representation, a tree works well to display a particular relationship and the distribution of the entities vis-24s that relationship. This tree shows who **is** on top and who is on the bottom of the chain of command. Some inferences can be made about prerogatives and responsibility, but only weakly since these inferences are based on pragmatic knowledge and not on knowledge that is stored in the structure of the classification itself. By knowing something about the domain, it is also possible to guess that GENERALS once were PRIVATES and thus bring "up the ladder" all of the experiences of going through the ranks, but this is not a formal requirement of the representation either, and may, in fact, be wrong.

Furthermore, a tree is "flatter" in its representation than is a hierarchy; there is less richness in the representation at each level because there is no inheritance or sharing of attributes. For example, there is no indication of the nature of LIEUTENANTS-their essence as it were-from their position in this classification. In a hierarchy, if we know a dog is a *mammal,* we know something about the mammalian attributes of the entity *dog.* What are the attributes of a LIEUTENANT that we can learn from the classification? Does **a** LIEUTENANT but has share attributes with a GENERAL less of them or different kinds? This type of information is not included explicitly. All we can know from this tree is that one rank commands the one lower in the pecking order.

Another kind of tree is one in which the entities are related by the partitive relationship. This means that each class is divided into its components, these components into subcomponents, and so on (see, for ex- ample, Figure 3).

In this example, SYRACUSE which in is *part* of ONONDAGA COUNTY, turn is *part of* NEWYORK, and so on. The partitive relationship (also known as *part/whole)* is a richer representation than the one shown in Figure 2. This is because the principle of inclusion allows more information to be shared. For instance, SYRACUSE COUNTY, is not *a kind of* ONONDAGA *so* what **is** true of the county *as a county* is not true for the city in it, but SYRACUSE is *part* Of ONONDAGA COUNTY and therefore inherits those attributes of the county that pertain to all units within it (e.g., location in New York State, climate, and so on). This relationship is so rich in representational power, in fact, that in many classification schemes there is no distinction made between the purely hierarchical and partitive.

NORTH AMERICA
- CANADA
  - ALBERTA
  - +
  - +
- MEXICO
  - +
  - CHIAPAS
  - +
- UNITED STATES
  - ALABAMA
  - +
  - +
  - NEW YORK
    - ONONDAGA COUNTY
      - SYRACUSE
  - +
  - +
  - WYOMING

Figure 3. Part/Whole Relationship.

relationships, and many people refer to both as "hierarchies." There is some psychological support for this, since both pure hierarchies and part/ whole classifications convey the notion of going from the more general and inclusive to the more specific or elemental.

Care must be taken, however, in making use of tree representations to ensure that the correct attributes are drawn upon in making inferences. This problem becomes clearer in another part/whole example (see Figure 4).

AUTOMOBILE
- BODY
- ENGINE BLOCK
  - PISTONS
  - VALVES
- INTERIOR
  - +
  - +

Figure 4. Part/Whole Relationship.

VALVES are *part* of the ENGINE BLOCK, but the nature of VALVES is distinct from the nature of PISTONS, and it would be incorrect to assume (despite their sibling position in the classification) that they share many attributes in common the way wolves and dogs do. In fact, VALVES and PISTONS are not similar entities at all. They share the attribute of being part of the ENGINE BLOCK, but that is only a partial explanation of what they are their essence. It would not be sufficient knowledge for most practical purposes the way knowing that dogs and wolves are closely related might prove useful. So, trees have the following formal requirements:

Complete and comprehensive information: Just like in a hierarchy, the entities that will be included in a tree must be decided in advance. First, it must be decided what will constitute an entity. Knowledge about the entities must be relatively complete in order to decide on the scope of the classification and the important criteria of distinction.

Systematic and predictable rules for distinction*:* The general structure of a tree is determined by the relationships among the entities. Part/ whole relationships might be appropriate for some knowledge, while other relationships (such as cause/effect; starting point/outcome; process/product; and so on) might be appropriate for other types of knowledge representation. These relationships should be ones that best reveal the knowledge of the domain-that is, the way in which all the entities interact with each other.

Citation order: In both hierarchies and trees it is important to decide the order in which rules of distinction will be invoked. The most important of these decisions is the “first cut” because this determines the shape and eventually the representational eloquence of the classification. If the first cut is a trivial one, the rest of the tree becomes awkward and does not reflect knowledge very well. For example, in the biological classification of animals (a hierarchy), the first cut is: *has a backbone/does not have a backbone (vertebrate/invertebrate).* While this cut produces a very skewed distribution in terms of numbers of species (there are many times more invertebrate species than vertebrates), the resulting classification proceeds smoothly down the sub- divisions and is able to cluster many attributes that “make sense” with respect to what we know about fundamental qualities of animals. In trees, the determination of an appropriate citation order is all the more important because trees are essentially descriptive, and the picture they present will depend on the first branching. For instance, in the AUTOMOBILE example presented above, it would be possible to make the first division BACKOF CAR/FRONT OF CAR, OF CAR/MIDDLE and proceed to decompose those sections into their component parts. But would this make sense? Would it present a reasonable division of an automobile’s components? Would it help us with knowledge about cars? Perhaps for someone in some context. There is no easy answer to what constitutes a meaningful division, and the decision often rests on consensual models or tradition.

Trees are useful knowledge representations for the following reasons:

Highlight/Display relationship of interest: This is the primary strength of a tree. It lays out the entities comprising a domain in a pattern of classes that highlights or makes evident the important or defining relationships among them.

Distance: A tree reveals the distances between entities (either physical distances or metaphorical ones). Thus one can determine that a COLONEL is “closer” to a GENERAL than is a PRIVATE at least along the dimensions of chain of command. If entities are components of the same super component, this means they are “closer” in space or in function.

Relative frequency of entities: This feature of trees is also shared by hierarchies. When entities cluster in large numbers under one classification label, this is frequently an opportunity for the creation or discovery of new rules for distinguishing among them. When a cluster is small and has only a few entities, these entities tend to be treated as if they were all the same. It may be neither feasible nor reasonable to make distinctions among them, and taking account of any differences may not support the enterprise. Once the cluster grows, however, and the number of entities reaches a critical mass, it might be used to further differentiate them. In such a case it is necessary to discover new knowledge that will suggest the best way of making these finer distinctions. Conversely when a category consistently has a member of one or just a few, it might signal the need for merging categories and rethinking the logic behind the division in the first place. In this case also, it is necessary to generate new knowledge in order to guide the merging or shifting of the orphan categories.

The use of trees as knowledge representations shares some of the same problems as does the use of hierarchies:

Rigidity: Because a tree is characterized by the relationships among entities and the citation order, the general shape of the tree-' its expressiveness as a knowledge representation-is determined a priori. This means that new entities can he added, if they fit into a place in the structure but, if the new entity or new knowledge does not fit well, the entire structure must be rethought and sometimes rebuilt.

One-way flow of information: In a hierarchy, information flows in two directions: vertically, between classes, super classes, and subclasses, and also laterally, between sibling classes (classes sharing the same super- class). In a tree, even if it is a part/whole representation, the information flows in a vertical direction up and down. Siblings in a class may in fact be entirely different types of objects. So there are rides for ***species*** but not for *differentia.* Many people assume that, since Syracuse is in New York State and New York City is also in New York State, that they are similar when in fact they only share the attribute of being in the same state and little else. Syracuse may be more like some other city in another state than it is like New York City. At any rate, the tree classification is not particularly good at representing multidirectional complex relationships.

Selective perspective: As with hierarchies, by emphasizing a certain relationship, a tree can mask, or fail to reveal, other equally interesting relationships. For instance, in the Army ranks example, the only relationship available to us is the "who commands whom" relationship. It does not touch upon the relationship of ranks when in combat, for instance, as opposed to the relationships among ranks of troops stationed at home. It does not show the distribution of men to women in the various ranks, or the distribution of ethnic or racial groups, and so on. It is completely silent on the classification of functional jobs in the Army (such as nurses, quartermasters, and so on). In other words, there are many other perspectives or lenses through which one could "know" the Army. The typology of ranks based on who commands whom is but one of them.

In summary, trees are useful for displaying information about entities and their relationships along one dimension of interest. They require fairly complete knowledge about a domain or at least about one aspect of a domain. A tree representation is good for displaying the relative placement of entities with respect to each other and their frequency at any node. On the other hand, trees are limited in how much they can represent, especially in terms of knowledge about entities within the same class. Furthermore, trees allow only partial inference.

Paradigms. A third classificatory structure is one in which entities are de- scribed by the intersection of two attributes at a time. The resulting matrix (or paradigm) reveals the presence or absence and the nature of the entity at the intersection (see Figure 5).

In this representation, we see two axes. The vertical has headings designating gender; the horizontal, types of kinship relationships. The cells represent the labels or names for this intersection of gender and kinship relationship. In this example, we have combined two such paradigms: one in English and one in Polish. You could imagine each one standing on its own but, for purposes of comparison, we have superimposed one on the other. Paradigms have the following formal requirements:

Two-way hierarchical relationship: Each cell entity (in this case the label signifying a kinship relationship) is related to both the vertical and the horizontal axis by a generic relationship. For instance, a *Mother is a Female,* and a *Mother* ***is*** *a Parent.* Across the row, all the entities are related to each other in being a subclass of the row header. So, *fathers, brothers, uncles, cousins,* and *fathers-in-law* are all *males.*

Kinship Relationship								
	Parent		*Sibling*		*Parent's Sibling*		*Parent's Sibling's Child*	
	Eng.	Pol.	Eng.	Pol.	Eng.	Pol.	Eng.	Pol.
Male	Father	Ojciec	Brother	Brat	Uncle	Stryj (father's side) Wujek (mother's side)	Cousin	Brat Stryjeczny (father's side) Brat Cioteczny (mother's side)
Female	Mother	Matka	Sister	Siostra	Aunt	Stryjenka (father's side) Ciocia (mother's side)	Cousin	Siostra Stryjeczna (father's side) Siostra Cioteczna (mother's side)

Figure 5. A Paradigm Displaying a Selection of KinshipTerms in English and Polish.

Entities in columns are related to each other by being in the same subclass. *So, uncles* and *aunts* are both *siblings of parents.* There is a shallow hierarchy running in both directions. However, entities are not related to each other in a generic relationship. Thus a *mother* is not a kind of *father* nor do they inherit properties from each other.

Axes represent two attributes of interest. Each axis represents one attribute that might serve to describe the entities in a meaningful way. In the example, the two axes represent *"the sex* of the person and *"the way the person is related."* The interesting feature of a paradigm is that it affords us a view of the entities classified along two dimensions at once.

Cells may be empty or may have more than one entity. Paradigms not only show us the intersection of two attributes, but also show us the presence, absence, and frequency of entities at these intersections.

So, how do we use this classificatory structure to represent and create knowledge?

Naming: Paradigms are frequently used in the study of terminology. As mentioned in a previous section, hierarchies enable the creation of strong definitions, but paradigms allow the study of patterns of naming. When people name things, they are creating an abstraction by incorporating a complex set of attributes under one label. Objects that are quite different in many ways but share defining attributes may still be given the same name. For instance, animals with a wide range of physical attributes are labeled *dog* if they share the defining attributes. Or, when we call something a *hamburger,* we may include under this rubric many slightly different kinds of sandwiches. They may have lettuce, or a slice of onion, or ketchup, or not; they may be small or large, but if they have a beef patty and a bun they are named *hamburger.* Now, if you add a slice of cheese, the name changes. Two hamburgers that are quite similar with respect to lettuce, tomato, on- ion, ketchup, and even sesame seeds, will still be named differently if one has a slice of American cheese. Naming will vary according to context, region, profession, and so on. So terminology indicates classificatory decisions, and paradigms serve as descriptive displays of terms as well as tools for analysis.

Distinction and lack of distinction: Paradigms can show the extent to which the intersecting criteria have distinct terms. In our example, we see that English has a single label for all the relationships displayed, while Polish has two terms each for uncles and aunts, and four terms for cousins, depending on whether they are related through the mother or father. So, in English, there is no distinction at all between cousins, and only a gender distinction between labels for your parents' sib- lings. Furthermore, besides being distinguished by side of the family, *cousins* in Polish are not completely distinguished from siblings and are given names that have the same root term as do *brothers* and *sisters.*

Patterns of similarity and difference: In terms of knowledge creation, paradigms often provide a heuristic tool for the discovery of regularity in the patterns of distinction. When distinctions are made in naming (that is, when people create different labels for concepts), we assume that the criterion for having made that distinction is important in some way. In this case, the distinction of relationship through either the mother's or father's side is important in Polish. Conversely, while English has distinct terms for cousins and siblings, Polish uses similar terms for both, distinguishing only by gender. The knowledge conveyed is that, even though each culture has a great deal of overlap in equivalent terminology, there are some subtle differences that may have historical or other explanations. In fact, it is interesting to note that the Polish distinction between aunts and uncles from different sides is fading. Did English once have such distinctions? Does this indicate the cultures are merging?

Empty cells: Empty cells in a paradigm provide an opportunity to investigate the reasons for the lack of a term. Does the absence of a term indicate the absence of a concept or does it indicate that the criteria chosen for the axes are not meaningful ones? Does every language, for instance, have the notion of a "kissing cousin?" Why is there no female equivalent of *misogynist?*

The limitations of paradigms as knowledge representation and discovery tools are as follows:

Requires knowledge of domain: The expressiveness of a paradigm relies on the felicitous choice of the attributes represented on the two axes. The fidelity of the picture that a paradigm reveals can be compromised if the dimensions are trivial and do not reflect fundamental concepts. In our example, the axes are chosen from concepts well established in the field of cultural anthropology: kinship expressed through blood and marriage, as well as distinctions made by gender. Paradigms that use dimensions guided by theory or a model usually do a better job of reflecting knowledge in the domain because they rest on a consensual framework of description. In other words, they are using a common vocabulary for communication. In fields where the fundamental relationships or concepts are not well understood, it is difficult to build a paradigm that reveals essential knowledge.

Limited perspective: While a well-chosen set of dimensions may pro- duce a valid description, it also produces a filter that limits the scope of what might be seen. So, in the example, kinship terms are ex- pressed using blood/marriage relationships and gender as dimensions of distinction. To us these distinctions seem self-explanatory and al- most universal, but in fact they are artifacts of our own cultural assumptions that we then impose on our observations of the world. Consider, for instance, that the dimensions do not address other family bonds, such as those that are based on strong affinity, legal adoption, and other socially invented forms of kinship. Nor do they allow for other cultural definitions of the entities themselves, such as alter- native views of what constitutes a *parent.* Thus, the paradigm presented in the example is a view through a particular lens. Another set of dimensions would present a different view and would most likely pro- duce different analytical outcomes.

Limited explanatory power: Because paradigms invoke dimensions only in pairs, they (like most classificatory structures) rarely produce a complete picture of a phenomenon. While paradigms do use the potentially rich representation of a hierarchical relationship vertically and horizontally, in a paradigm this relationship is shallow-only one deep-and therefore not very complex. For this reason, paradigms are essentially descriptive. They help clarify; they may *suggest* patterns and anomalies, but these patterns are not inherent to the structure and must be interpreted by the person using the paradigm.

In summary, paradigms are good tools for discovery. They reveal the presence or absence of names for entities defined by pairs of attributes. They can be used for comparison and for the display of patterns and anomalies with respect to the variety and distribution of terms. Paradigms are heuristic in that they present a clear view that can then be analyzed and interpreted. Like most classificatory structures, paradigms require knowledge of the domain or some guiding principles in order to make a good choice of dimensions and, like most classificatory structures, paradigms are usually partial and biased represantations.

Faceted Analysis. Faceted classifications are not really a different representational structure but rather a different approach to the classification process. The notion of facets rests on the belief that there is more than one way to view the world and that even those classifications that are viewed as stable are in fact provisional and dynamic. The challenge is to build classifications that are flexible and can accommodate new phenomena.

Faceted classification has its roots in the works of S.R. Ranganathan, an Indian scholar, who posited that any complex entity could be viewed from a number of perspectives or *facets.* He suggests that these fundamental categories are Personality, Matter, Energy, Space, and Time (Ranganathan, 1967). Over the years, Ranganathan’s facets have been reinterpreted in many contexts, but it is surprising how well they have weathered the test of time. They have been used to classify objects as disparate as computer software (for reuse), patents, books, and art objects (Kwasnik, 1992a).

Not all faceted classifications use Ranganathan’s prescribed fundamental categories, but what they do have in common is the process of analysis.

Period/Style	*Place*	*Process*	*Material*	*Object*
19th Century	Japanese	raku	ceramic	vase
Arts & Crafts	American		oak	desk

Figure 6. A Faceted Analysis of Artifacts.

Figure **6** shows a possible solution to the classification of material culture which, in its diversity, defies easy description and categorization. For purposes of demonstration, this is a simplified version of the one used by the *Art and Architecture Thesaurus.* For any given artifact, there are many possible ways of representing it, let alone the “knowledge” that enabled its production or its value. The faceted approach follows these steps:

Choose facets: Decide, in advance, on the important criteria for description. These form the facets or fundamental categories. In this case we have Period, Place, Process, Material, and Object, following closely on what Ranganathan suggested.

Develop facets*:* Each facet can be developed/expanded using its own logic and warrant and its own classificatory structure. For example, the Period facet can be developed as a timeline; the Materials facet can be a hierarchy; the Place facet a part/whole tree, and so on.

Analyze entities using the facets: In analyzing an entity, choose descriptors from the appropriate facets to form a string, as shown above. Thus, the classification string for object 1 is “19‘”Century Japanese raku ceramic vase.” The string for object 2 is “Arts & Crafts American oak desk.” It is important to note that the process is not one of division (as in a hierarchy) where the entities are subdivided into ever more specifically differentiated categories. It is not a process of *decomposition* either (as in a part/whole tree), in which the entities are broken down into component parts, each part different from the whole. Instead, the process of *analysis* is to view the object from all its angles-same object but seen from different perspectives. So, in the example, the vase can be seen from the point of view of its period, the place in which it was made, the material and processes, and so on.

Develop citation order. In organizing the classified objects, choose a primary facet that will determine the main attribute and a citation order for the other facets. This step is not required and applies only in those situations where a physical (rather than a purely intellectual) organization is desired.

The development of a faceted approach has been a great boon to classification because it meshes well with our modern sensibilities about how the world is organized. Specifically, it is a useful tool because it:

Does not require complete knowledge: In building a faceted scheme, it is not necessary to know either the full extent of the entities to be accommodated by the scheme nor the full extent of the relationships among the facets. It is thus particularly useful in new and emerging fields or in fields that are changing.

Hospitable. When a classification is hospitable it means it can accommodate new entities smoothly. In a faceted scheme, if the fundamental categories are sound, new entities can be described and added. This is particularly important in the classification of objects such as cultural artifacts, where we have no way of predicting the things that will be produced by the human imagination. If an artifact produced 100years from now could be described by the fundamental categories of period, place, material, process, and object, then the classification scheme will still be robust.

Flexibility. *Since a faceted scheme describes each object by a number of independent* attributes, these attributes can be invoked in an endlessly flexible way in a sort of Leg0 approach. "Let me see all the iron objects made in 17^{th} century Scotland." "OK, now all the copper objects." "OK, now iron objects in Italy.. .". This flexibility can be used to discover new and interesting associations. The approach is called *post-coordination* and means that attributes can be mixed and matched at the time of retrieval. It is in contrast to the *pre-coordinated* categories that are a requirement of most hierarchies in which the rules for class inclusion are invoked at the time the entity is classified and stay fixed from there on.

Expressiveness. A faceted approach can be more expressive because each facet is free to incorporate the vocabulary and structure that best suits the knowledge represented by that facet.

Does not require a strong theory. Since a faceted classification does not have an overall structure, it does not have to have a "theoretical glue" to hold it together and to guide the rules for association and distinction. It can be constructed ad hoc so long as the fundamental categories function well.

Can accommodate a variety of theoretical structures and models. A faceted approach makes it possible to represent a variety of perspectives as well. For instance, in facet analyzing a piece of literature, one facet may reflect a particular model of genres, another a model of languages, and so on. In a traditional hierarchy, it may be extremely difficult or impossible to blend the two, while a faceted scheme allows them to co- exist.

Multiple perspectives. One of the most useful features of a faceted approach is that it allows entities to be viewed from a variety of perspectives-a feature that **is** lacking in hierarchies and trees. In a faceted analysis, it is possible to describe a *dog* as an *animal,* as a *pet,* as *food,* as a *commodity*, and ad infinitum, *so* long as the fundamental categories have been established with which to do this.

While the flexibility and pragmatic appeal of faceted classifications have made this a popular approach, there are some limitations in terms of knowledge representation and creation:

Difficulty of establishing appropriate facets. The strength of a faceted classification lies in the fundamental categories, which should express the important attributes of the entities being classified. Without knowledge of the domain and of the potential users, this is often difficult to do. While it is possible to flexibly add entities, it is not a simple matter to add fundamental facets once the general classification is established.

Lack of relationships among facets. Most faceted classifications do not do a good job of connecting the various facets in any meaningful way. Each facet functions as a separate kingdom, as it were, without much guidance as to how to put the parts together. For example, to facet analyze motion pictures by *genre, country, director, film process,* and *so* on, we would still have no insight as to the meaningful relationships of, say, a particular country and the popular film genre there or of a particular film process arid the genres it supports.

In terms of theorizing and model building, the faceted classification serves as a useful and multidimensional description but does not explicitly connect this description in an explanatory framework.

Difficulty of visualization. A hierarchy or a tree, and especially a paradigm, can be visually displayed in such a way that the entities and their relationships are made evident. This is difficult to do for a faceted classification, especially if each facet is structured using a different internal logic. As a result, faceted schemes can only be viewed along one or two dimensions at a time, even though a more complex representation is actually incorporated into the descriptive strings. Thus it is difficult to see a vase in the context of other vases, of other Japanese artifacts, of other clay objects, of other raku objects, and so on, *all at the same time.*

Nevertheless, faceted schemes continue to flourish because we recognize that they allow at least some systematic way of viewing the world without the necessity for a mature and stable internal framework in which to view it. Information technology has promise for new ways of enabling multidimensional visualization and for developing computer-assisted ways of discovering patterns and anomalies that can possibly lead to new knowledge.

CLASSIFICATION AND KNOWLEDGE

There are many ways in which classification schemes and knowledge interact. Sometimes the interaction is so harmonious that the two remain linked for a long time. Sometimes knowledge changes and the classification must also change or knowledge changes and the classification is no longer adequate to the task. Sometimes the classification itself generates new knowledge. The following discussion is representative of ways in which knowledge and classifications mutually interact.

Changing, Explanatory Frameworks

The Periodic Table of Elements, attributed to Mendeleyev, is an example of a classification scheme that has endured through several explanatory Frameworks. When the Periodic Table was first proposed, there was already a body of knowledge about individual elements-i.e., facts and observations, including the knowledge of atomic weight. It was observed that elements could be arranged in a systematic order according to atomic weight, and this would show a periodic change of properties. This early Periodic Table proved to be a very useful tool, leading to the discovery of new elements and a new understanding of already known elements. In terms of theory, the Table "so divided its subject matter that it [could] enter into many and important true propositions about the subject matter..." (Kaplan, 1963, p. 50). The rule for determining one element from another was atomic weight, a basically descriptive criterion, but it was not until the discovery of atomic theory that the periodicity of the table was fully understood. This theory *explained* (rather than merely described) the underlying principles behind the regularity and pattern of the classified entities. With this new explanation, many new properties could be inferred. The table became a predictive tool for as yet undiscovered elements as well as an explanatory tool and a very fruitful descriptive tool.

It reflected well what was already known about elements and pointed to new knowledge (such as the common characteristics of inert gases). What is interesting is that the original Periodic Table did not have to undergo fundamental changes in structure even though a new explanatory framework was discovered.

Changes in Perspective

Technological advances in measuring and viewing instruments have had a profound influence on classifications. This is because new instruments reveal new knowledge that does not always fit neatly into existing knowledge representation structures. Such instruments have included carbon dating, the electron microscope, DNA testing, remote sensing, and so on. For instance, clouds were traditionally very simply classified by shape and by height from the horizon. This classification was developed when we could only see clouds from our perspective standing on the earth. Now we can measure the moisture, temperature, particulate matter, and charge of a cloud. Moreover, clouds can be observed from a satellite thereby observing global patterns. We know about the typical life cycle of a cloud-how clouds change shape and identity. This new understanding of clouds has a profound impact on weather forecasting, navigation, and other fields of knowledge, yet the traditional classification is robust as a simple and clear form of communication. It remains very popular and co- exists with new classifications.

Sometimes the new way of observation yields a new classification. For instance, gems used to be classified on a scale of hardness (with diamonds at one end and chalk at the other) and also by color. These attributes were visible to the naked eye. Once it was possible to view gems through a microscope, it was then necessary, or at least more useful, to alter this classification to include knowledge about crystalline structure.

Changing Entities

Sometimes, though, changes in the way we can observe have led to fundamental changes not only in the classificatory structures but also in the nature of the entities themselves. For instance, a complex problem exists in trying to coordinate traditional ways of describing natural habitats with new ways of observation and measurement. Formerly, scientists spoke in terms of rainfall, temperature, growth forms, dominant life forms, and so on. These attributes were described and classified according to whatever model predominated or was accepted and resulted in such constructs as *deserts*, *tropical rainforests*, and *so* on. But today the mapping of natural habitats is done by remote sensing data that measure a different kind of unit such as reflectance, texture, density, spatial patterns, slope, Leaf Area Index, and so on (Muchoney, 1994)-that is, each of these parameters does not necessarily correspond to something (an entity) that we can unambiguously call "a tree" or a "a camel." Remote sensing data do not require a semantically coherent entity to be the cause of measurement. All that matters is that measurements can be taken and this measurement can then be combined, clustered, and analyzed with other measurements to yield "types." Put another way, the resulting entities are not semantically meaningful in the same way as traditional names. The problem arises in communication about the habitats and also in theorizing about them-i.e., in making sense of the phenomena.

Another example of changing entities within a domain has occurred in the classification of musical instruments. Instruments are classified basically by material and the method of producing sound (striking, blowing, bowing, strumming, and so on). This classification did quite well with a few slight adjustments here and there for hybrid instruments, but it hit a real snag with the introduction of electronic instruments such as synthesizers. The problem lies not so much in squeezing the newcomer categories into the old scheme but rather in how well the old criteria fit the new entities as meaningful guides for association and discrimination. This situation is a classic example of shifts that occur only occasionally at first but might eventually lead to a complete overhaul of the classification or perhaps to the creation of two parallel classifications.

Classifications Built When There is No Consensus

Many classifications must be built when there is no generally accepted theory or model on which to construct or to define the entities. For ex- ample, the classification of mental disorders is mandated economically by insurance companies and legal requirements. Social institutions require that we be able to determine who is legally sane, who must be confined to care, and who will be reimbursed for services. This classification is there- fore built on factors that are not based on any particular theory of mental illness or mental processes but rather on readily observable symptoms and behaviors. It is therefore only a moderately good descriptive tool and an almost useless tool for understanding fundamental processes from any theoretically coherent perspective. Thus, the classification fails to act as a heuristic device by generating provocative questions or providing interesting insights.

Classifications Where There are No uniform or Stable Entities

Scientists have long been struggling with a classification of smells. While we are able to construct useful classifications of colors based on what we know about the physics of light, the psychology and aesthetics of color and the human ability to perceive and use colors, we have not had the same success with smells. We are forced to refer to smells using analogies: *fruity, citrus, green, floral, putrid*, and *so* on. One of the problems is that there is no "unit" of smell, no building block that could then be classed and differentiated in some systematic way. This lack of a classification has led to reliance on the essentially subjective artistry of individuals with respect to the identification, blending, and general understanding of smell. This does not mean that we do not know how smell works in terms of perception. Nor does it mean that we do not understand the powerful nature of smell in human life, but the fact remains that we have no good way of talking about smells as smells.

A similar problem arises in the classification of viruses. This is not because viruses do not exist as entities but, rather, because they change and are never in a form unambiguous enough to be pinned down by a clear classification.

The Intersection of Theory and Economic Interests

Classifications are never created in a political or social vacuum. Everyone is familiar with the old Department of Agriculture classification of food groups: meats, dairy and eggs, grains, fruits and vegetables, and fats. From a classification point of view, the dairy and eggs category always seemed a rather odd one in that it is not clear on what basis dairy products and eggs go together, nor along which dimension they are distinguished from the other categories. It is not by source (animals) nor by nutritional component (protein). Furthermore, the classification does not indicate the relative importance of each group. “At least one from each” was the slogan. As it turns out, the classification was the result of a strong lobby by the dairy industry. It is amazing how well established this classification became and how long it persisted. The new classification is really quite elegant. It builds a pyramid with grains on the bottom and fats on the top. Dairy and eggs now share a level with meats. This new classification reflects modern nutritional science much more coherently. It not only classes the foods according to some understandable criteria, but also indicates, by the narrowing pyramidal shape, the relative amount of foods from each of the groups that should be consumed. This classification may, as it turns out some day, reflect faulty scientific knowledge, but at least it reflects it with fidelity and clarity.

Keys and Other “Thin” Classifications

A key is a classification that is built using an easily identifiable, but not necessarily theoretical, set of criteria. One example is a field guide to flowers. In such a guide, flowers are classed first by their petal color. Petal color is a characteristic that is easy to identify but is trivial when compared to more fundamentally meaningful attributes of a flower such as plant structure or reproductive mechanisms. The petal color yields relatively little fundamental knowledge of the flower, but petals are an easy way to narrow the field of possibilities, especially for the novice. A key, there- fore, allows easy entry into a deeper classification. However, consider the classification of baseball in the Dewey Decimal Classification in Figure 7.

796	Athletic and outdoor sports and games
796.3	Ball games
796.35	Ball driven by club, mallet, bat
796.357	***Baseball***

Figure 7. Dewey Decimal Classification of Baseball.

This scheme positions baseball, along with other ball games in which a ball is hit with a mallet, club, or bat, right next to field hockey and croquet as well as polo. This is not an inaccurate classification. There are no factual errors, as it were, but by using the ball and the bat as the defining criteria, the classification is reduced to a very thin representation of what baseball is. It does not address any of the team aspects; the cultural aspects; or the aesthetic, athletic, economic, strategic, or spiritual aspects of baseball as a sport. Why? Because this is difficult to do, and there is no consensus really of what these attributes might be or how they might be expressed. There is no generally accepted theory of games or sports, but we can all readily agree that yes, indeed, in baseball you hit a ball with a bat. In this case, the classification is structured like a key but does not lead to a deeper theoretical representation of baseball in all its complexity.

CONCLUSION

Classification is a way of seeing. Phenomena of interest are represented in a context of relationships that, at their best, function as theories by providing description, explanation, prediction, heuristics, and the generation of new questions. Classifications can be complex or simple, loaded with information or rather stingy in what they can reveal. They can reflect knowledge elegantly and parsimoniously, or they can obfuscate and hinder understanding. Some classifications enable flexible manipulation of knowledge for the purposes of discovery; some are rigid and brittle, barely able to stand up under the weight of new knowledge. It is useful to understand the properties of various classification structures so we can exploit their strengths and work around the weaknesses. In the future, classification will be enhanced by new methods of revealing patterns, associations, and structures of knowledge, and by new ways of visualizing them.

Notes:

Some of the ideas in this article were first described in a paper presented at the Third ASIS SIG/CR Workshop on Classification Research, Pittsburgh, PA, 1992 (Kwasnik, l992b).

Adapted from an example in Aitchison, Jean and Gilchrist, Alan (eds.). (1987). *Thesaurus construction*, 2“ed. London, England: Aslib, p. *80*.

REFERENCES

Ackrill, J. L. (Trans.). (1963). *Aristotle's "Categories "and "De Interpretatione": Translated with notes.* Oxford, England: Oxford University Press.

Bronowski, J. (3978). *The origins of knowledge and imagination.* New Haven, CT Yale University Press.

Davies, R. (1989). The creation of new knowledge by information retrieval and classification. *Journal of Documentation, 45(4),*273-301.

Kwasnik, B. H. (1992a). The legacy of facet analysis. In R. N. Sharma (Ed.), *S.R.* ***Ranganathan and the West*** (pp. *98-3* 11). New Delhi, India: Stelling.

Kwasnik, B. H. (1992b). The role of classification structures in reflecting and building theory. In R. Fidei, B. H. Kwasnik, & P. J. Smith (Eds.), *Advances in classification research, vol. 3* (Proceedings of the 3rd ASIS SIG/CR Classification Research Workshop) (pp: 63-81), Medford, NJ: Learned Information, for the American Society for Information Science.

Muchoney, D. M. (1996). Relationships and divergence of vegetation and mapping classifications. In R. Fidel, C. Beghtol, B. H. Kwasnik, & P. J. Smith (Eds.), *Advances in classification research, vol. 5* (Proceedings of the 5"ASIS SIC/CR Classification Research Workshop). Medford, NJ: Learned Information, for the American Society for Information Science.

Olson, H.**A.** (1998). Mapping beyond Dewey's boundaries: Constructing classificatory space for marginalized knowledge domains. *Library trends, 47(2),*233-254. Ranganathan, S. R. (1967). *Prologomena to library classification, 3rd* ed. Bombay: Asia Publishing House.