

A REPRESENTATION FOR NATURAL CATEGORY SYSTEMS

Sandra L. Peters and Stuart C. Shapiro
Department of Computer Science
State University of New York at Buffalo
Buffalo, NY 14260
peters@buffalo.csnet, shapiro@buffalo.csnet

ABSTRACT

Most AI systems model and represent natural concepts and categories using uniform taxonomies, in which no level in the taxonomy is distinguished. We present a representation of natural taxonomies based on the theory that human category systems are non-uniform. That is, not all levels of abstraction are equally important or useful; there is a basic level which forms the core of a taxonomy. Empirical evidence for this theory is discussed, as are the linguistic and processing implications of this theory for an artificial intelligence/natural language processing system. We present our implementation of this theory in SNePS, a semantic network processing system which includes an ATN parser-generator, demonstrating how this design allows our system to model human performance in the natural language generation of the most appropriate category name for an object. The internal structure of categories is also discussed, and a representation for natural concepts using a prototype model is presented and discussed.

1. INTRODUCTION.

Knowledge-base systems typically model and represent natural concepts and categories using *uniform* inheritance networks [Quillian 1967, 1968, 1969; Collins & Quillian 1970; Fahlman 1979] or frame systems [Brachman 1983; Brachman & Schmolze 1984]. We will present a representation of natural taxonomies based on the theory that human category systems are non-uniform, i.e., that not all levels of abstraction are equally important or useful. This theory is supported by a substantial body of empirical evidence from the fields of psychology, anthropology, and linguistics [Rosch et al. 1976, 1978; Mervis & Rosch 1981; Berlin 1978; C. H. Brown et al. 1976; Tversky 1978; Hunn 1976; Cantor et al. 1979; Smith & Medin 1981]. We will discuss some of the evidence for this theory, as well as some of the linguistic and processing implications of this theory for an AI system modeling human cognitive behavior.

The need for a non-uniform representation will be extended as we consider the internal structure of natural concepts. We will present and discuss some of the empirical evidence that supports the use of a prototype model for these concepts, and present a representation using this model. We will argue, however, that certain levels of abstraction exhibit more of a prototypicality structure than others, and that, therefore, distinct representations are again needed in modeling the internal structure of concepts at different levels of abstraction in a taxonomy.

Our implementation uses the SNePS semantic network processing system which includes an ATN parser-generator [Shapiro, 1978, 1979, 1982, 1986].

2. THEORY THE VERTICAL DIMENSION OF CATEGORY SYSTEMS - A BASIC LEVEL.

Our representation is based on the following principles of human categorization set forth by Eleanor Rosch. Categories within taxonomies are structured such that there is one level of abstraction at which the most basic category cuts can be made. This level of abstraction forms the "core" [Berlin 1978, p. 24] of a taxonomy, and is called the basic level. Basic categories are: (1) those which carry the most information; (2) those whose members have the most attributes in common; and (3) the categories

most differentiated from one another. Basic level categories are, in fact, disjoint. Chair, car, and dog are examples of basic level objects.

Levels of a taxonomy above the basic level are called superordinate categories (e.g., furniture, vehicle, mammal). Fewer attributes are shared among members of superordinate categories, i.e., there is less category resemblance. Categories below the basic level are called subordinate categories (e.g., kitchen chair, station wagon, collie). Subordinate categories contain many attributes which overlap with those of other subordinate categories, i.e., there is less contrast between categories across a subordinate level.

2.1. Empirical Evidence.

The following summarizes some of Rosch's empirical evidence supporting the existence of a basic level which forms the core of a taxonomy. [Rosch et al. 1976, 1978; Mervis & Rosch 1981].

2.1.1. Attributes of Objects.

When subjects were asked to list attributes of basic, superordinate, and subordinate level objects, very few attributes were listed for superordinate categories, a great number of attributes were listed for basic categories, and an insignificant number of additional attributes were listed for subordinate level categories. This result supports the theory that the basic level is the most inclusive or general level at which the objects of a category possess a large number of attributes in common. Attributes appear to be clustered at the basic level.

2.1.2. Object Recognition.

Experiments using averaged shapes, obtained by superimposing outlines of objects to form normalized shapes, showed that the basic level is the most inclusive level at which the averaged shape of an object can be recognized. That is, basic objects (e.g., chairs, dogs) were the most general objects that could be identified from these shapes; superordinate objects (e.g., furniture, animals) could not be identified from averaged shapes. This suggests that basic level objects, are the most inclusive categories for which a concrete mental image of the category as a whole can be formed. We can form an image of a cat or dog which reflects the average members of the class, however, we cannot form an image of a mammal that reflects the appearance of the class as a whole.

2.1.3. Object Names Categorization.

Studies of picture verification have demonstrated that objects are first recognized as members of their basic level category. When subjects were shown pictures of objects, the basic level name was the name chosen for an object. With additional processing time, subjects were able to categorize objects at their subordinate and superordinate levels. Thus, subjects knew the subordinate and superordinate names of objects, but categorized objects first at the basic level. Rosch further states that basic level objects are the first categorizations made during perception of the environment, as well as the categories most named, and most necessary in language.

2.1.4. Development of Categories.

Basic level objects are not only the first categories learned by children, they appear to be formed differently from categories at other levels. That is, basic categories are not learned explicitly by acquiring a definition or deductive rule, but rather are learned implicitly by exposure to multiple instances of the category; i.e., they are formed inductively. This is often called the acquisition of types through ostensive definitions [Jackendoff 1983]. Categories subordinate and superordinate to this level are often formed by the acquisition of a deductive rule [Berlin 1978]. For example, the concept mammal might be learned in terms of a rule which lists attributes such as: warm-blooded; body usually covered with hair; female gives milk to young.




Figure 1a




Figure 1b

2.1.5. Summary of Empirical Evidence.

Thus, recent categorization research provides a great deal of empirical evidence supporting the importance of basic level categories in a taxonomy, and the non-uniformity of human category systems. Basic level categories are the first categories developed, they are formed differently than non-basic categories, they are the most used and useful categories, and therefore, they must be distinguished from non-basic categories in some way.

3. REPRESENTATION AND USE OF CATEGORIES IN AN AI/NLP SYSTEM.

If an artificial intelligence/natural language processing (AI/NLP) system modeling human category systems must be able to distinguish basic level categories from non-basic categories, an important issue to be considered is how and where to make the distinction. Basic level objects are used in two kinds of categorization: "ordinary" categorization, i.e., the classification of an individual in a class, and generic categorization, i.e., categorization involving two classes or types. It seems clear that since basic level categories are formed early in life, they are formed via ordinary categorization. The teaching of these names is limited to the presentation of examples and counter-examples. Thus, a child may learn the basic level name 'dog', as someone points to Rover and says 'dog'. Therefore, our system makes the distinction between basic and non-basic levels in the representations for ordinary categorization, i.e., in the individual/class relations.

Figure 1a shows the case frame used for ordinary categorization of a basic level object. Here **m1** represents the proposition that the individual represented by **i** is a member of the basic level category represented by **j**. Figure 1b shows the representation for "Rover is a dog". [See Shapiro & Rapaport 1986 for the syntax and semantics of other constructs.]

Since non-basic categories are formed later than basic categories, and are formed in the course of the investigation of underlying principles rather than ostensive features, we use a slightly more complex case frame to represent membership in a non-basic level category. Thus, in Figure 2a **m1** represents the proposition that the individual represented by **i** is a member of the non-basic category represented by **j**. The representation of "Rover is a mammal" is shown in Figure 2b.




Figure 2a




Figure 2b




Figure 3a




Figure 3b




Figure 3c

These case frames in SNePS are the built-in syntactic structures of our modeled minds. The use of the *member/class* case frame reflects the basic or primitive nature of categorization in basic categories, whereas the use of the *arg1/rel/arg2* case frame treats membership in non-basic categories as an ordinary binary relation. Thus, our system distinguishes two cases of ordinary categorization: one representation is used when the class membership involves a basic level category, another representation, when the class membership involves a non-basic category.

In addition to this ordinary categorization, a system must, of course, be able to represent generic categorization, i.e., class/class relations, such as "Dogs are mammals". These relations are represented using a subclass/superclass case frame. See Figure 3a. Here *m1* represents the proposition that the class of *i*'s are a subclass of the class of *j*'s. The representations of "Dogs are mammals" and "Collies are dogs" are shown in Figures 3b and 3c respectively.

Thus, we build a traditional uniform type hierarchy of class/class relations. We see no reason to distinguish any relations in the hierarchy, since we find no evidence that generic categorization sentences such as "collies are dogs", "dogs are mammals", "mammals are vertebrates", require different underlying representations. (It is noteworthy that there are no class/class relations between two basic level categories.)

Since an ability to form abstract concepts is required for generic categorization, this categorization occurs at a later stage of development than does ordinary categorization of basic level objects. Therefore, the type hierarchy, which is formed after basic level concepts are formed, is not the appropriate place to make the distinction between basic and non-basic categories. In summary, a single representation is used for class/class relations, but two distinct representations are used for individual/class relations.

KRL-0 [Bobrow & Winograd 1977a, 1977b, 1979] is, to our knowledge, the only other AI system to distinguish basic and non-basic levels in the representation of taxonomies. KRL-0 used *units* to represent both classes and individuals. Three distinct levels of abstraction were used in the representation of classes or types in *units*: a basic level, an abstract level, and a specialization level. Bobrow and Winograd stated that they did not, however, find an appropriate way to use these unit categorization levels for classes, and removed unit categorization from KRL-1 [Bobrow & Winograd 1979, p. 41]. Although not precisely specified in their papers, Bobrow and Winograd appear to have made distinctions among the levels of abstraction in the type hierarchy of frames only, not in the individual/class relationships. We could not find any evidence that distinctions were made in the units representing individuals [Bobrow & Winograd 1977a p. 23].

3.1. INHERITANCE AND LINGUISTIC IMPLICATIONS.

3.1.1. Inheritance.

One of the organizational principles to which most semantic networks and frame systems adhere is that of storing properties in the hierarchy at the place covering the maximal subset of nodes sharing them. This is an efficient organizational scheme in which properties do not have to be replicated at

different places in the network, for they are inherited by nodes below the ones in which they are stored. This principle fits in well with the theory of cognitive economy, for one can gain a great deal of information from a category system organized in this way, while conserving resources.

Categorization research studies, however, do not support this principle of organization. As stated above, properties appear to be clustered at the basic level, not at the level covering the maximal subset of nodes. This means that there is not a great deal of inheritance of properties taking place in the type hierarchy. Instead most inheritance occurs at the individual level, i.e., from the basic level category to the individual. Thus, Rover inherits attributes from the basic level category dog.

3.1.2. Linguistic Implications.

Perhaps the most dramatic enhancement to our system resulting from our distinguishing basic and non-basic level categories is our ability to model human performance by choosing the most appropriate category name for an object. Systems using uniform taxonomies have to make arbitrary word choice decisions. For example, the NIGEL generator [Sondheimer et al. 1986] generates as specific a term as possible. However, we know from human categorization research that in the absence of a specific context that would lead one to use a non-basic level name for an object, the basic level name should be used.

The dialog shown in Figure 4 illustrates our system's ability to model human performance in this respect. Since the basic level name is the most useful and most used name, the most appropriate answer to the question "What did Lucy pet?" is not the specialization "collie" or the superordinate level name "mammal, but the basic level name "dog".

The additional dialog shown in Figure 5 demonstrates that the basic level name is chosen regardless of the order in which categories are mentioned, and also demonstrates the natural language input of classificational information.

4. THEORY - THE INTERNAL STRUCTURE OF CATEGORIES.

Although categories have been viewed traditionally as concepts established by necessary and sufficient criteria, and many AI systems have modeled natural concepts in this way [Brachman 1983; Brachman & Schmolze 1984; Fahlman 1979], recent categorization research on naturally occurring concepts does not support this view [Rosch 1976, 1978; Mervis & Rosch 1981; Smith & Medin, 1981]. Rosch has suggested that another way to achieve separateness and clarity of categories is by "conceiving of each category in terms of its clear cases" [Rosch 1978, p. 36], that is, in terms of prototypes. Thus, all members of a category are not equally representative of the category, but, rather, some exemplars are more representative than others. Prototypical category members are those members which have the most attributes in common with other members of the category and the fewest attributes in common with members of contrasting categories.

4.1. Empirical Evidence Supporting Prototype Theory. A large body of empirical evidence exists which supports the theory that categories are conceived of in terms of prototypes, and that, therefore, category members vary in representativeness [Rosch 1978; Mervis & Rosch 1981; Smith & Medin 1981]. A brief overview of some supporting evidence follows.

4.1.1. Production of Exemplars. When subjects are asked to list exemplars of superordinate categories, they list the most representative exemplars. Similarly, when subjects are asked to sketch an exemplar for a particular category, they sketch the most representative exemplars [Mervis & Rosch, 1981].

4.1.2. Categorization: Speed of Processing. Reaction times studies show that the verification of a robin as a bird is performed faster than the verification of an ostrich or penguin as a bird. Thus, subjects are able to verify category membership faster for those objects rated as being prototypical, i.e., for representative exemplars, than for those rated as non-prototypical [Mervis & Rosch, 1981]. In addition,

atn parser initialization

: Lucy petted a yellow animal
I understand that Lucy petted a yellow animal

: The animal was a dog
I understand that the yellow animal is a dog

: The dog was a collie
I understand that the yellow dog is a collie

: What did Lucy pet
Lucy petted a yellow dog

: ^end
(end atn parser)

Figure 4

atn parser initialization

: Mary petted a dog
I understand that Mary petted a dog

: The dog is a mammal
I understand that the dog is a mammal

: The dog was a labrador
I understand that the dog is a labrador

: What did Mary pet
Mary petted a dog

: Jane petted a manx
I understand that Jane petted a manx

: The manx is a cat
I understand that the manx is a cat

: A cat is a mammal
I understand that cats are mammals

: Mammals are animals
I understand that mammals are animals

: Who petted an animal
Mary petted a dog
and
Jane petted a cat

: ^end
(end atn parser)

Figure 5

when a prime, i.e., the prior mention of the category name, is provided, the response time to verify category membership of representative exemplars decreases. Priming, however, increases the response time necessary to verify the membership of non-representative exemplars [Mervis & Rosch, 1981].

4.1.3. Learning and Development of Categories.

Category membership is established first for those exemplars that are most representative of the category; membership for non-representative exemplars vacillates for some time. Thus, the formation of category prototypes is related to the initial formation of categories: the most representative members of categories are the ones first established as category members. In addition, categories are learned more easily if initial exposure is confined to representative exemplars [Mervis & Rosch 1981].

4.1.4. Categorization: Indeterminacy of Category Membership. Category boundaries are not well defined. Many experiments have revealed that subjects disagree concerning the category to which poor exemplars belong [Berlin & Kay 1969; Labov 1973; McCloskey & Glucksberg 1978]. Both Sokal [1974] and Rosch & Mervis [1975] have demonstrated that poor exemplars contain attributes that overlap with those of contrast categories.

4.1.5. Perception of Typicality Differences. There is general agreement among subjects when they are asked to rate how good an example is of its category, or to choose the exemplar most representative of its category. [Rosch 1978; Mervis & Rosch 1981].

4.2. Summary of Empirical Evidence. If natural categories are determined by necessary and sufficient criteria, then category members should be equivalent, i.e., any member of the category should be as good an example of the concept as any other one. The empirical evidence outlined above does not support this view. Reaction time studies showing the speed of categorization, studies of the learning and development of categories, studies of typicality ratings, and studies involving the production of exemplars, demonstrate that all members are not equally good examples of their category, but rather that some members of a category are more representative than others.

We are not proposing that a prototype model is an appropriate one for all concepts, but are considering only natural concepts in this paper. Some classes of entities may indeed be determined by necessary and sufficient criteria.

4.3. Representation.

As Rosch points out, use of a prototype model does not specify a representation of categories, it merely constrains our choice of representation [Rosch 1978]. Different theories of semantic memory can accommodate this view. However, it seems clear that a representation based on the classical model, i.e., on necessary and sufficient criteria, cannot accommodate the evidence discussed above.

4.3.1. Additional Constraints.

We agree with the proposal [Rosch 1978] that any representation for a natural concept should satisfy the requirements of (1) mirroring the structure in the perceived world and (2) cognitive economy. Cognitive economy dictates that a representation for the concept *dog* should include only the essential information for the categorization of novel instances of this concept, i.e., information about how dogs resemble one another and differ from other concepts. In addition, we believe that cognitive economy further dictates that we mentally represent categories in terms of an *abstraction* which is an amalgamation of the most salient and most modal features of category members. An alternative to this featural model is the mental image model of prototypes. However, as Johnson-Laird [1983] points out, images are highly specific. Thus, we do not form an image of trees in general or of fruit in general, rather we form an image of a specific tree or a specific fruit. We cannot form an image that is isomorphic to the class as a whole for these categories. The image model would require mental images for each of the clear-cut cases or typical category members, rather than a composite of the features of these typical members, and yet seems no more useful in categorizing novel instances. Thus, the

featural model is a more economical one. The featural model is also able to capture functional attributes which cannot easily be captured in a mental image or holistic model.

Our representation makes use of one further constraint: our belief that types or categories are non-projectable or non-referring, i.e., types do not exist. As Jackendoff [1983] points out, we cannot point to a type, but only to instances of types. We do not consider a prototype to be an additional member of its category, but rather consider it to be an abstraction: a list of features. Therefore, although our representation for the prototypical *dog* may contain the feature *four-leggedness*, we do not concern ourselves with structural information such as this prototypical dog's typical left front leg. We are only interested in the abstract feature or property *four-leggedness*, for its facilitation of categorization. This constraint allows us to avoid some of the inheritance problems with which Fahlman's NETL [1979] must deal, since he must concern himself that an individual elephant does not inherit the typical elephant's typical left front leg. (See also Shapiro [1980] for a review of NETL.)

4.4. Implementation of Prototype Model in SNePS.

Our implementation represents a category as a collection of abstract features or properties. A separate proposition is used to capture each abstract feature. Thus, the category *bird* is represented in terms of the most salient and modal features of the members of the concept *bird*, e.g., *flies*, *feathered*, *winged*, *has a beak*, *sings*, etc. Each of these properties is represented with a default generalization, e.g., *flies* is represented with a default generalization that may be paraphrased as: For all x if x is a bird, then presumably x flies. Figure 6 shows a partial representation for the prototype *bird*, using two of the features mentioned above. Thus, a prototype in our system consists of a bundle of default generalizations, all of which share the same antecedent, x is a bird, in our example, and the same variable.

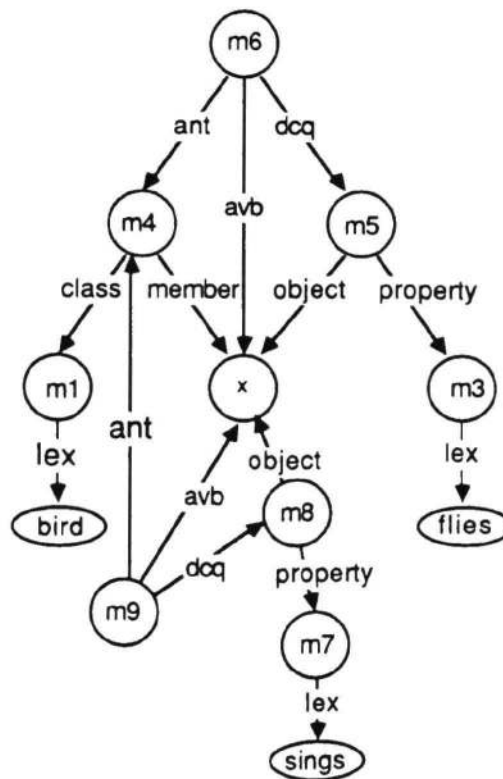


Figure 6

4.5. Effect of the Level of Abstraction on the Representation.

While superordinate level categories are formed in the course of the investigation of underlying principles, basic level categories are formed by exposure to multiple instances, on the basis of ostensive features. Therefore, it seems that basic level level categories should exhibit more of a prototypicality structure than superordinate level categories. The basic level is also the level of abstraction at which attributes are clustered; very few attributes cluster at higher levels in the taxonomy. Thus, it seems appropriate to use distinct representations for the internal representation of basic and non-basic level categories: a prototype representation for basic level categories and perhaps, a mixed representation that combines deductive rules which are true universals, e.g., mammals are warm-blooded, with an exemplar component, that is, with a listing of prominent exemplars, for superordinate level concepts.

4.6. Understanding Generic Sentences.

This work grew out of an interest in understanding generic sentences, i.e., sentences containing generic concepts. We are now considering whether the prototype model is of use in understanding these sentences when they involve basic level or subordinate level concepts. One of the problems these sentences pose for an AI/NLU system is that a multiplicity of quantifiers can be posited for sentences of this type. Consider the following sentences:

- (1) Cardinals are red. (The mature males are red.)
- (2) Mosquitos carry malaria. (Very few actually carry malaria.)
- (3) Dogs are four-legged. (Most are four-legged.)
- (4) Shamrocks are green. (Most are green.)
- (5) Men are mortal. (Mortality is a necessary criterion in this case.)

It may be more appropriate to examine the properties or features expressed in terms of saliency, rather than to try to determine the sense of quantification in these sentences. Features which we consider to be especially salient, such as *transmits dangerous disease*, may be part of the abstraction that form the prototypical mosquito, regardless of the sense of quantification in the sentence or utterance.

4.7. PROCESSING IMPLICATIONS.

The non-uniformity of human category systems also has implications for a processing model for categorization. Category research has established that objects can be identified as members of their basic level category more rapidly than as members of their superordinate or subordinate categories. A possible processing model for our implementation, compatible with Rosch's empirical evidence and the current general processing assumptions about categorization involving featural models [Smith & Medin 1981] such as ours is the following. An object is first identified or recognized as a member of its basic class, since properties or attributes are clustered at the basic level. Because of this bundling of attributes at the basic level, this processing involving feature matching can be performed quickly. Categorization of an object as a member of its subordinate classes requires additional processing time, because additional features must be matched, some of which are much less salient than the features for categorizing an object at the basic level. Categorization of an object as a member of its superordinate classes requires inferencing using the type hierarchy. We use path-based inference to accomplish this [Shapiro 1986, 1978]. Performing this inferencing, of course, requires additional processing time.

5. CONCLUSIONS.

We have incorporated principles of categorization derived from several years of research in our AI/NLP system. Principles of categorization, taken seriously, affect the design and representation of taxonomies for natural concepts. We distinguish one level, the basic level, as the core of our

taxonomies, using a representation for membership in basic categories distinct from that used for membership in non-basic categories. This representation allows our system to model human performance in the generation of appropriate names for objects: when there is no context effect, the basic level name is used.

The use of a prototype representation and storing of attributes at the basic level will allow us to model human performance in categorization tasks involving basic and non-basic level objects.

REFERENCES

- (1) Ashcraft, M. H. (1978), "Property Norms for Typical and Atypical Items from 17 Categories: A Description and Discussion," *Memory and Cognition*, vol. 6, pp. 227-232.
- (2) Berlin, B. (1978), "Ethnobiological Classification," *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates) pp. 9-27.
- (3) Berlin, B., Kay, P. (1969), *Basic Color Terms: Their Universality and Evolution* (Berkeley: Univ. Calif. Press).
- (4) Bobrow, D. G., Winograd, T. (1977a), "Experience with KRL-0, One Cycle of a Knowledge Representation Language," *IJCAI-77*, vol. 1, pp. 213-222.
- (5) Bobrow, D. G., Winograd, T. (1977b), "An Overview of KRL, a Knowledge Representation Language," *Cognitive Science*, vol. 1:1, pp. 3-46.
- (6) Bobrow, D. G., Winograd, T. (1979), "KRL: Another Perspective," *Cognitive Science*, vol. 3, pp. 29-42.
- (7) Brachman, R. J. (1979), "On the Epistemological Status of Semantic Networks," *Associative Networks: Representation and Use of Knowledge by Computers* (New York: Academic Press), Edited by N. V. Findler, pp. 3-50.
- (8) Brachman, R. J., Fikes R. E., Levesque, H. J. (1983), "KRYPTON: A Functional Approach to Knowledge Representation," *IEEE Computer*, vol. 16, pp. 67-73.
- (9) Brachman, R. J., Schmolze, J., (1984), "An Overview of the KL-ONE Knowledge Representation System," Fairchild Technical Report Number 655, (Palo Alto, CA: Fairchild Laboratory for Artificial Intelligence Research), September 1984.
- (10) Brachman, R. J., (1985), "'I Lied About the Trees' Or, Defaults and Definitions in Knowledge Representation," *The AI Magazine*, (Fall 1985), pp. 80-93.
- (11) Brown, C. H., Kolar, J., Torrey, B. J., Truong-Quang, T., Volkman, P. (1976), "Some General Principles of Biological and Non-Biological Folk Classification," *Am. Ethnol.*, vol. 3, pp. 73-85.
- (12) Brown, R. (1958), "How Shall a Thing Be Called?," *Psychol. Rev.*, vol. 65, pp. 14-21.
- (13) Cantor, N., Mischel, W. (1979), "Prototypes in Person Perception," *Adv. Exp. Soc. Psychol.*, vol. 12, pp. 3-52.
- (14) Carlson, G. N. (1982), "Generic Terms and Generic Sentences," *Journal of Philosophical Logic*, vol. 11, pp. 145-181.
- (15) Collins, A. M., Quillian M. R., (1970), "Does Category Size Affect Categorization Time?," *Journal of Verbal Learning and Verbal Behavior*, vol. 9, pp. 432-438.
- (16) Delgrande, J. P., (1986), "A Propositional Logic for Natural Kinds," *Proc. Canadian Conf. on AI-86*, pp. 44-48.
- (17) Fahlman, S. E., (1979), *NETL: A System for Representing and Using Real-World Knowledge* (Cambridge, MA: MIT Press).
- (18) Gaschnig, J., (1979), "Preliminary Performance Analysis of the PROSPECTOR Consultant System for Mineral Exploration," *IJCAI*, vol. 6, pp. 308-310.

- (19) Hunn, E., (1976), "Toward a Perceptual Model of Folk Biological Classification," *Am. Ethnol.*, vol. 3, pp. 508-524.
- (20) Jackendoff, R., (1983), *Semantics and Cognition* (Cambridge, MA: The MIT Press).
- (21) Johnson-Laird, P.N., (1983), *Mental Models* (Cambridge, MA: Harvard U. Press).
- (22) Labov, W., (1973), "The Boundaries of Words and Their Meanings," *New Ways of Analyzing Variations in English* (Washington: Georgetown Univ. Press) Editors: C. J. Bailey, R. Shuy, pp. 340-373.
- (23) Kitts, D. B., Kitts, D. J., (1979), "Biological Species as Natural Kinds," *Philosophy of Science*, vol. 46, pp. 613-622.
- (24) Lehnert, W., Wilks, Y., (1979), "A Critical Perspective on KRL," *Cognitive Science*, vol. 3, pp. 1-28.
- (25) McCloskey, M. E., Glucksberg, S., (1978), "Natural Categories: Well Defined or Fuzzy Sets?" *Mem. Cognit.*, vol. 6, pp. 462-472.
- (26) Melle, W. van, Shortliffe, E. H., Buchanan, B. G., (1981), "EMYCIN: A Domain-independent System that Aids in Constructing Knowledge-Based Consultation Programs," *Machine Intelligence, Infotech State of the Art Report*, vol. 9:3.
- (27) Mervis, C. B., Rosch, E., (1981), "Categorization of Natural Objects," *Ann. Rev. Psychol.*, vol. 32, pp. 89-115.
- (28) Quillian, M. R., (1967), "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities," *Behavioral Science*, vol. 12, pp. 410-430.
- (29) Quillian, M. R., (1968), "Semantic Memory," *Semantic Information Processing* (Cambridge, MA: MIT Press), Editor, M. Minsky.
- (30) Quillian, M. R., (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language," *Communications of the ACM*, vol. 12, pp. 459-476.
- (31) Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., Boyes-Braem, P., (1976), "Basic Objects in Natural Categories," *Cognitive Psychology*, vol. 8, pp. 382-439.
- (32) Rosch, E., Lloyd, B. B., (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- (33) Shapiro, S. C., (1978), "Path-Based and Node-Based Inference in Semantic Networks," *TINLAP-2: Theoretical Issues in Natural Language Processing* (New York: ACM), Editor, D. Waltz, pp. 219-225.
- (34) Shapiro, S. C. (1979), "The SNePS Semantic Network Processing System," *Associative Networks* (New York: Academic Press), Editor, N.V. Findler, pp. 179-203.
- (35) Shapiro, S. C. (1980), "Book Review: NETL: A System for Representing and Using Real-World Knowledge," *Amer. Journal of Comp. Ling.*, vol. 6, pp. 183-186.
- (36) Shapiro, S. C., (1981), "COCCI: A Deductive Semantic Network Program for Solving Microbiology Unknowns," Technical Report Number 173, (Buffalo: SUNY Buffalo Dept. of Computer Science).
- (37) Shapiro, S. C., (1982), "Generalized Augmented Transition Networks Grammars for Generation from Semantic Networks," *Amer. Journal of Comp. Ling.*, vol. 8, pp. 12-25.
- (38) Shapiro, S. C., Rapaport, W. J., (1986), "SNePS Considered as a Fully Intensional Propositional Semantic Network," *Proc. AAAI-86*, vol. 1, pp. 278-283.
- (39) Shapiro, S. C., (1986), "Symmetric Relations, Intensional Individuals, and Variable Binding," *Proc. of the IEEE*, vol. 74, no. 10, pp. 1354-1363.
- (40) Smith, E. E., Medin, D. L., (1981), *Categories and Concepts* (Cambridge, MA: Harvard University Press).

- (41) Sondheimer, N. K., Nebel, B. (1986), "A Logical-Form and Knowledge-Base Design for Natural Language Generation," *AAAI-86*, vol. 1, pp. 612-618.
- (42) Tversky, A., Gati, I., (1978), "Studies of Similarity," *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 81-95.
- (43) Wegner, P., (1986), "The Object-Oriented Classification Paradigm," Technical Report No. CS-86-11, (Providence, RI: Brown University Dept. of Computer Science).
- (44) Woods, W., (1978), "Semantics and Quantification in Natural Language Question Answering," *Advances in Computers* (New York: Academic Press), Editor, M. C. Youvits, vol. 17, pp. 2-87.