IJCAI 87

Proceedings of the Tenth International Joint Conference on Artificial Intelligence

Volume 1

August 23–28, 1987
sponsored by
International Joint Conferences on Artificial Intelligence, Inc. (IJCAII)
cosponsored and hosted by
Associazione Italiana per l'Informatica ed il Calcolo Automatico-A.I.C.A.
(Gruppo di Lavoro di Intelligenza Artificiale GLIA)
in cooperation with
Politecnico di Milano, Italy

Copyright © 1987, International Joint Conferences
on Artificial Intelligence, Inc.
All rights reserved.

Edited by
John McDermott

Distributed by
Morgan Kaufmann Publishers, Inc.
95 First Street
Los Altos, California 94022

ISBN 0-934613-43-5

Printed in the Netherlands
through Interprint, San Francisco.

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

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. There is a basic level which forms the core of a taxonomy; both higher and lower levels of abstraction are less important and less useful. 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. Among these implications are: (1) when there is no context effect, basic level names should be used; (2) systems should identify objects as members of their basic level categories more rapidly than as members of their superordinate or subordinate categories. 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 ability of our system to acquire classificational information from natural language sentences is also demonstrated.

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., 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.

This work is part of a larger, ongoing research effort concerned with problems in the understanding of natural language sentences containing generic terms. We will demonstrate some of the current abilities of our system to acquire generic concepts from natural language sentences, and to use these generic concepts in answering questions and making categorization judgements. This 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 also 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.

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. The case frame used for this form of categorization of a basic level object is:

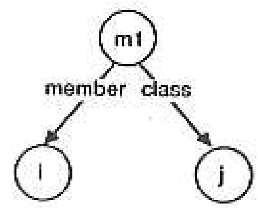

Figure 1

Here m1 represents the proposition that the individual represented by i is a member of the basic level category represented by j. "Rover is a dog" is represented as follows:

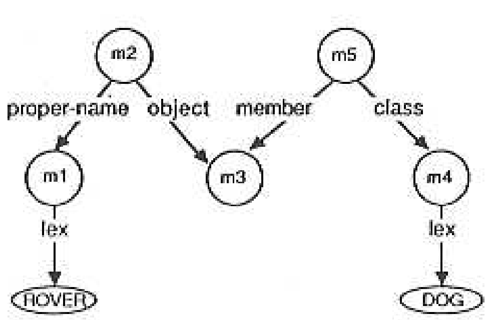

Figure 2

[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.

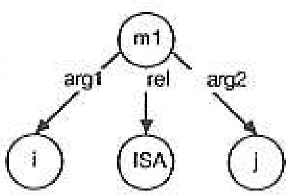

Figure 3

Here m1 represents the proposition that the individual represented by i is a member of the non-basic category represented by j. "Rover is a mammal" is represented as follows:

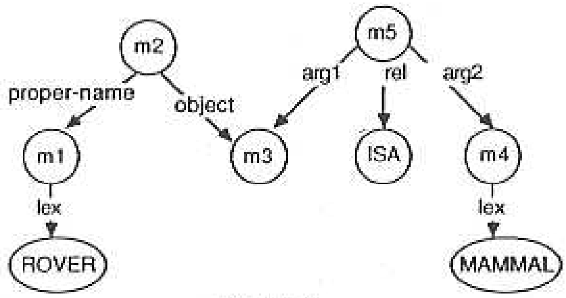

Figure 4

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.

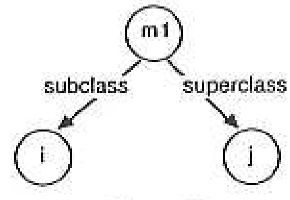

Figure 5

Here m1 represents the proposition that the class of i's are a subclass of the class of j's. "Dogs are mammals" is represented as follows:

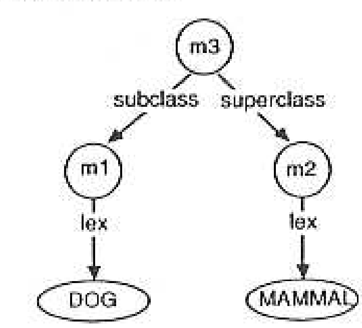

Figure 6

Likewise, "collies are dogs" is represented as follows:

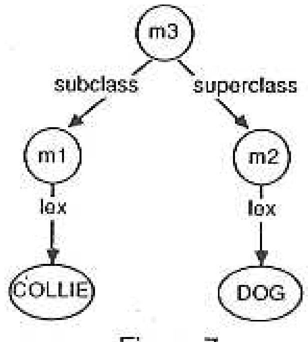

Figure 7

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].

4. INHERITANCE AND LINGUISTIC IMPLICATIONS.

4.1. 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.

4.2. 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.

Figure 8 shows a dialog with our system illustrating our ability to model human performance in this respect:

```
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 8

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 dialog in Figure 9 demonstrates that the basic level name is chosen regardless of the order in which categories are mentioned.

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 9

Thus, Figures 8 and 9 show that word choice decisions are not made arbitrarily. Our system does not simply choose the most or least specific name of an object, or the category name mentioned either first or most recently in the dialog. Rather, the most appropriate name for an object, its basic level name, is used.

Figure 9 also shows the use of both ordinary categorization information and generic categorization information. The first three sentences show ordinary categorization, i.e., the multiple classifications of an individual as *dog, mammal,* and *collie.* Figure 10 shows part of the network built following the input of these sentences. Node m9 represents the individual classified as a *dog, mammal,* and *collie.* The basic level name, *dog,* is chosen to answer the question "What did Mary pet".

Figure 11 shows part of the network constructed from the input of the last five sentences in the dialog. Node m24 is the individual classified both as a *cat* and *manx.* The last group of sentences in Figure 9 also includes two generic categorizations: "A cat is a mammal" and "Mammals are animals". A type hierarchy is constructed from this input. Answering the question "Who petted an animal" requires inferencing using the type hierarchy. Our system has inheritance rules which make this inferencing possible. [See Shapiro 1978, and Shapiro & Rapaport 1986 for examples of these rules.]

5. PROCESSING IMPLICATIONS.

The non-uniformity of human category systems also has implications for a processing model for categorization. We would like to use our system to model classification problem solving. This form of problem solving is the basis for many expert systems, e.g., PROSPECTOR [Gauschnig 1980], EMYCIN [van Melle 1979], and COCCI [Shapiro 1981] are knowlege-base systems that specialize in forms of classification problems.

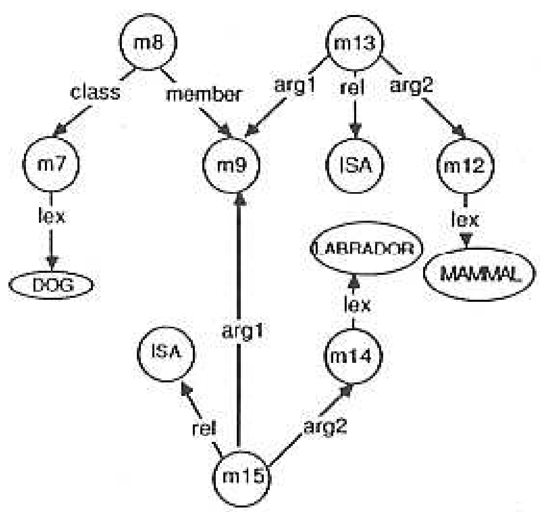

Figure 10

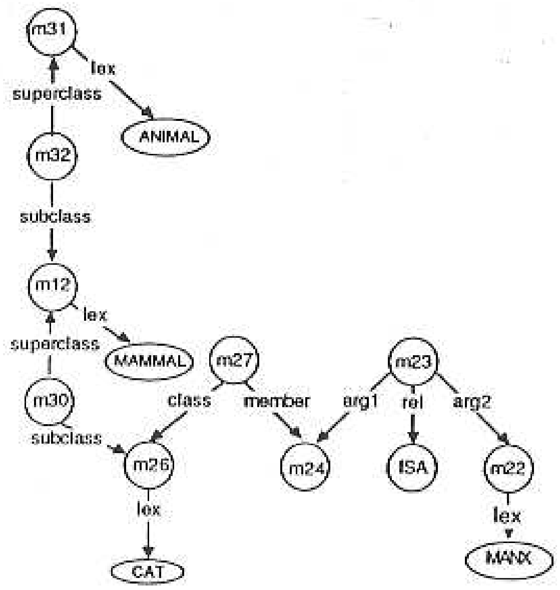

Figure 11

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 & Rapaport 1986, Shapiro 1978]. Performing this inferencing, of course, requires additional processing time.

We are also interested in modeling the effect of expertise on the classification system, since Rosch [1976] found that expertise affects which level of abstraction is considered to be the basic level, as well as the amount of information stored at the basic and subordinate levels. For example, an airplane mechanic participating in her studies did not treat *airplane* as a basic level category, but further differentiated airplanes to form basic level categories. His list of attributes for types of airplanes was much more lengthy than those of other subjects, and he used attributes ignored by others. His visual view of airplanes also differed from those of other subjects, since his canonical view of airplanes was of the undersides and the engines, rather than of the top and side images. Although it seems clear that the effects of expertise will be confined to small, specific parts of the taxonomy, the effects of expertise on the classification system need to be studied further. We believe that our system is flexible enough to accomodate the effects of expertise on the organization of the system.

6. CONCLUSIONS.

We have incorporated principles of categorization derived from several years of research in our AI/NLP system. 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 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 distinct representations and storing of attributes at the basic level also will allow us to model the additional processing time necessary to categorize an object at a non-basic level.

REFERENCES

(1) Ashcraft, M. H. (1978), "Property Norms for Typical and Atypical Items from 17 Categories: A Description and 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) Bobrow, D. G., Winograd, T. (1977a), " Experience with KRL-0, One Cycle of a Knowledge Representation Language," *IJCAI-77*, vol. 1, pp. 213-222.

(4) Bobrow, D. G., Winograd, T. (1977b), "An Overview of KRL, a Knowledge Representation Language," *Cognitive Science*, vol. 1:1, pp. 3-46.

(5) Bobrow, D. G., Winograd, T. (1979), "KRL: Another Perspective," *Cognitive Science*, vol. 3, pp. 29-42.

(6) 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.

(7) Brachman, R. J., Fikes R. E., Levesque, H. J. (1983), "KRYPTON: A Functional Approach to Knowledge Representation," *IEEE Computer*, vol. 16, pp. 67-73.

(8) 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.

(9) Brachman, R. J., (1985), ""I Lied About the Trees" Or, Defaults and Definitions in Knowledge Representation," *The AI Magazine*, (Fall 1985), pp. 80-93.

(10) 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.

(11) Brown, R. (1958), "How Shall a Thing Be Called?," *Psychol. Rev.*, vol. 65, pp. 14-21.

(12) Cantor, N., Mischel, W. (1979), "Prototypes in Person Perception," *Adv. Exp. Soc. Psychol.*, vol. 12, pp. 3-52.

(13) Carlson, G. N. (1982), "Generic Terms and Generic Sentences," *Journal of Philosophical Logic*, vol. 11, pp. 145-181.

(14) 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.

(15) Delgrande, J. P., (1986), "A Propositional Logic for Natural Kinds," *Proc. Canadian Conf. on AI-86*, pp. 44-48.

(16) Fahlman, S. E., (1979), *NETL: A System for Representing and Using Real-World Knowledge* (Cambridge, MA: MIT Press).

(17) Gaschnig, J., (1979), "Preliminary Performance Analysis of the PROSPECTOR Consultant System for Mineral Exploration," *IJCAI*, vol. 6, pp. 308-310.

(18) Hunn, E., (1976), "Toward a Perceptual Model of Folk Biological Classification," *Am. Ethnol.*, vol. 3, pp. 508-524.

(19) Jackendoff, R., (1983), *Semantics and Cognition* (Cambridge, MA: The MIT Press).

(20) Kitts, D. B., Kitts, D. J., (1979), "Biological Species as Natural Kinds," *Philosophy of Science*, vol. 46, pp. 613-622.

(21) Lehnert, W., Wilks, Y., (1979), "A Critical Perspective on KRL," *Cognitive Science*, vol. 3, pp. 1-28.

(22) Melle, W. van, Shortliffe, E. H., Buchanan, B. G., (1981), "EMYCIN: A Domain-independent System that Aids in Constructing Knowlege-Based Consultation Programs," *Machine Intelligence, Infotech State of the Art Report*, vol. 9:3.

(23) Mervis, C. B., Rosch, E., (1981), "Categorization of Natural Objects," *Ann. Rev. Psychol.*, vol. 32, pp. 89-115.

(24) Quillian, M. R., (1967), "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities," *Behavioral Science*, vol. 12, pp. 410-430.

(25) Quillian, M. R., (1968), "Semantic Memory," *Semantic Information Processing* (Cambridge, MA: MIT Press), Editor, M. Minsky.

(26) Quillian, M. R., (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language," *Communications of the ACM*, vol. 12, pp. 459-476.

(27) 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.

(28) Rosch, E., Lloyd, B. B., (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates).

(29) 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.

(30) Shapiro, S. C. (1979), "The SNePS Semantic Network Processing System," *Associative Networks* (New York: Academic Press), Editor, N.V. Findler, pp. 179-203.

(31) 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).

(32) Shapiro, S. C., (1982), "Generalized Augmented Transition Networks Grammars for Generation from Semantic Networks," *Amer. Journal of Comp. Ling.*, vol. 8, pp. 12-25.

(33) Shapiro, S. C., Rapaport, W. J., (1986), "SNePS Considered as a Fully Intensional Propositional Semantic Network," *Proc. AAAI-86*, vol. 1, pp. 278-283.

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

(35) Sondheimer, N. K., Nebel, B. (1986), "A Logical-Form and Knowledge-Base Design for Natural Language Generation," *AAAI-86*, vol. 1, pp. 612-618.

(36) Tversky, A., Gati, I., (1978), "Studies of Similarity," *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 81-95.

(37) Wegner, P., (1986), "The Object-Oriented Classification Paradigm," Technical Report No. CS-86-11, (Providence, RI: Brown University Dept. of Computer Science).

(38) 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.

(7) Brachman, R. J., Fikes R. E., Levesque, H. J. (1983), "KRYPTON: A Functional Approach to Knowledge Representation," *IEEE Computer*, vol. 16, pp. 67-73.

(8) 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.

(9) Brachman, R. J., (1985), ""I Lied About the Trees" Or, Defaults and Definitions in Knowledge Representation," *The AI Magazine*, (Fall 1985), pp. 80-93.

(10) 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.

(11) Brown, R. (1958), "How Shall a Thing Be Called?," *Psychol. Rev.*, vol. 65, pp. 14-21.

(12) Cantor, N., Mischel, W. (1979), "Prototypes in Person Perception," *Adv. Exp. Soc. Psychol.*, vol. 12, pp. 3-52.

(13) Carlson, G. N. (1982), "Generic Terms and Generic Sentences," *Journal of Philosophical Logic*, vol. 11, pp. 145-181.

(14) 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.

(15) Delgrande, J. P., (1986), "A Propositional Logic for Natural Kinds," *Proc. Canadian Conf. on AI-86*, pp. 44-48.

(16) Fahlman, S. E., (1979), *NETL: A System for Representing and Using Real-World Knowledge* (Cambridge, MA: MIT Press).

(17) Gaschnig, J., (1979), "Preliminary Performance Analysis of the PROSPECTOR Consultant System for Mineral Exploration," *IJCAI*, vol. 6, pp. 308-310.

(18) Hunn, E., (1976), "Toward a Perceptual Model of Folk Biological Classification," *Am. Ethnol.*, vol. 3, pp. 508-524.

(19) Jackendoff, R., (1983), *Semantics and Cognition* (Cambridge, MA: The MIT Press).

(20) Kitts, D. B., Kitts, D. J., (1979), "Biological Species as Natural Kinds," *Philosophy of Science*, vol. 46, pp. 613-622.

(21) Lehnert, W., Wilks, Y., (1979), "A Critical Perspective on KRL," *Cognitive Science*, vol. 3, pp. 1-28.

(22) Melle, W. van, Shortliffe, E. H., Buchanan, B. G., (1981), "EMYCIN: A Domain-independent System that Aids in Constructing Knowlege-Based Consultation Programs," *Machine Intelligence, Infotech State of the Art Report*, vol. 9:3.

(23) Mervis, C. B., Rosch, E., (1981), "Categorization of Natural Objects," *Ann. Rev. Psychol.*, vol. 32, pp. 89-115.

(24) Quillian, M. R., (1967), "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities," *Behavioral Science*, vol. 12, pp. 410-430.

(25) Quillian, M. R., (1968), "Semantic Memory," *Semantic Information Processing* (Cambridge, MA: MIT Press), Editor, M. Minsky.

(26) Quillian, M. R., (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language," *Communications of the ACM*, vol. 12, pp. 459-476.

(27) 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.

(28) Rosch, E., Lloyd, B. B., (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates).

(29) 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.

(30) Shapiro, S. C. (1979), "The SNePS Semantic Network Processing System," *Associative Networks* (New York: Academic Press), Editor, N.V. Findler, pp. 179-203.

(31) 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).

(32) Shapiro, S. C., (1982), "Generalized Augmented Transition Networks Grammars for Generation from Semantic Networks," *Amer. Journal of Comp. Ling.*, vol. 8, pp. 12-25.

(33) Shapiro, S. C., Rapaport, W. J., (1986), "SNePS Considered as a Fully Intensional Propositional Semantic Network," *Proc. AAAI-86*, vol. 1, pp. 278-283.

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

(35) Sondheimer, N. K., Nebel, B. (1986), "A Logical-Form and Knowledge-Base Design for Natural Language Generation," *AAAI-86*, vol. 1, pp. 612-618.

(36) Tversky, A., Gati, I., (1978), "Studies of Similarity," *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 81-95.

(37) Wegner, P., (1986), "The Object-Oriented Classification Paradigm," Technical Report No. CS-86-11, (Providence, RI: Brown University Dept. of Computer Science).

(38) 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.