The Tenth Annual Conference of the Cognitive Science Society

August 17-19, 1988

Montreal, Quebec, Canada

Program of the Tenth Annual Conference of the Cognitive Science Society

17-19 August 1988

Montreal, Quebec, Canada

LEA LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
1988 Hillsdale, New Jersey Hove and London

Distributed by

Lawrence Erlbaum Associates, Inc.
365 Broadway
Hillsdale, New Jersey 07642

ISBN 0-8058-0436-6

Printed in the United States of America

FLEXIBLE NATURAL LANGUAGE PROCESSING AND ROSCHIAN CATEGORY THEORY

Sandra L. Peters, Stuart C. Shapiro, and William J. Rapaport

Department of Computer Science
State University of New York at Buffalo

1. INTRODUCTION.

Artificial intelligence systems typically hand-craft large amounts of knowledge in complex, static, high-level knowledge structures. These systems generally work well in very limited domains, but are simply too rigid to support natural language understanding in general. This is due in part to the fact that artificial intelligence, natural language processing (AI/NLP) systems have not taken seriously the principles of categorization first set forth in the seminal work of Rosch [1976, 1978, 1981], and later extended by researchers including Barsalou [1987], Murphy & Medin [1985], Lakoff [1987], and Neisser [1987]. Thus, AI/NLP systems have very shallow representations of generic concepts and categories, employing either static, simple, featural models of concepts based on necessary and sufficient criteria in uniform taxonomies where no level is distinguished [e.g., KLONE, Brachman 1983], or passive data structures with slots and explicit default values, such as frames, schemata, and scripts [e.g., KRL, Bobrow & Winograd 1977ab; NETL, Fahlman 1979]. We have previously shown that systems that use the former representations are unable to model human category systems [Peters & Shapiro 1987ab]. In this paper, we discuss the inadequacy of systems based on the latter types of representations, arguing that frames, schemata, and scripts lack the flexibility, generality, and adaptability necessary for representing generic concepts in memory. We present alternative "active" representations, in which frames or schemata do not reside in semantic memory, but rather are constructed as needed from a less organized semantic memory. Their construction can, therefore, be influenced by the current task and context.

In addition, our processing and representations are based on a Roschian model of categories, i.e., on (1) a recognition of the unique nature of basic level categories within natural category systems, and (2) prototype theory. We will discuss some of the current research that supports our representations and processing, and show that our system's performance is enhanced by taking these principles of categorization seriously. Our implementation uses the SNePS knowledge representation and reasoning system, including a generalized ATN parser-generator [Shapiro 1979, 1982; Shapiro & Rapaport 1987]. We present a detailed example that shows the use of our representations and processing strategies in the task of discourse comprehension. In particular, our example will concern implicit focusing in natural language comprehension, i.e., the implicit activation of thematic associates and salient attributes of concepts during discourse comprehension.

1.1. Roschian Model: Basic Level Primacy. Psychology, linguistics, and anthropology have produced a variety of measures of perception, behavior, and communication showing a convergence of cognitive tasks at the basic level. Not all levels of a taxonomy are equally used and useful: for taxonomies of common objects and organisms, the basic level, the level of *table* and *bird*, is the most informative and useful [Rosch et al. 1976; Berlin 1978; Tversky & Hemenway 1984]. Our knowledge is organized at this level (i.e., most attributes of category members are stored at the basic level) and visual imagery is particularly strong for basic level concepts. The informativeness of the basic level originates from the amount of knowledge stored at this level and the rich perceptual component of basic level categories. Basic level concepts trigger many reflex inferences; i.e., they routinely activate many satellite concepts. Where informativeness is the greatest, so too is inferential power.

We have previously presented a representation for natural category systems that uses distinct representations for basic and non-basic level concepts, and discussed empirical evidence supporting basic level primacy and the need for AI/NLP systems to recognize the uniqueness and importance of the basic level [Peters & Shapiro 1987ab]. The work described in this paper builds on, and presupposes that earlier work.

1.2. Roschian Model: Prototype Theory. Although categories have been viewed traditionally as concepts established by necessary and sufficient criteria, and many AI systems continue to model natural concepts in this way [e.g., Brachman 1983], recent categorization research does not support this view [Rosch 1976, 1978; Mervis & Rosch 1981; Murphy & Medin 1985; Barsalou 1987; Lakoff 1987; Neisser 1987]. Rosch has suggested that another way to achieve the separateness and clarity of categories is by conceiving of each category in terms of its clear cases, i.e., prototypes. Categories possess graded structure: the members of a category vary in how typical they are of their category [Rosch, et al. 1976; Barsalou 1985; Lakoff 1987].

1.3. Extending the Roschian Model. Although Rosch [1978] cautioned that the existence of prototype effects merely indicates that prototypes must have some place in theories of representation of categories, many have misinterpreted her work, construing prototypes as a complete theory of representation for categories and generic concepts. In particular, many AI researchers [e.g., Minsky 1975; Bobrow & Winograd 1977b; Fahlman 1979; Schank 1977] have considered prototypes to constitute a representation for generic concepts, and have employed passive data structures with slots and explicit default values to directly represent these concepts in memory. Thus, they have used frames, scripts, and schemata as a basis for models of cognitive processing, i.e., using them to represent generic objects, situations, events, sequences of events, actions, and sequences of actions. Many cognitive scientists have recently pointed out that passive representations such as frames are too rigid, lacking the flexibility, generality, and richness required in cognitive processing [e.g., Barsalou 1987; Kintsch 1987; Lakoff 1987; Neisser 1987]. Both Schank [1985] and Minsky [1986] have also recently proposed the need for more flexible, dynamic representations. Empirical evidence supporting the need for active representations of generic concepts and categories will be discussed in the following sections.

2. EMPIRICAL EVIDENCE.

2.1. The Instability of Graded Structure of Categories. Barsalou has demonstrated that the graded structure of a category is unstable, varying greatly across contexts. He concludes that "different concepts temporarily represent the same category in working memory on different occasions"[1987, p. 101]. A category's graded structure shifts as a function of (1) the linguistic context and (2) the point of view from which it is perceived [Barsalou 1987]. For example, when *animal* is processed in the context of farm, *cows, horses, goats*, and *pigs* are more typical than *bears, lions, elephants*, and *giraffes*. This situation is reversed when *animal* is processed in the context of zoo.

2.2. The Instability of Generic Concepts. Barsalou [1987] proposes that there are no invariant concepts in long-term memory; rather, long-term memory contains large amounts of highly interrelated knowledge that is used to construct concepts in working memory. Generic concepts are not retrieved intact from long-term memory when needed: they are constructed from long-term memory for a particular task in a particular context. In Barsalou's view, generic concepts have both context-independent and context-dependent information associated them. Barsalou [1982] has presented evidence that shows that context-independent properties are automatically activated by presented concepts, while context-dependent properties are activated only in particular situations or contexts. Thus, the representations for a generic concept in working memory may vary widely across contexts.

2.3. The Inflexibility of Scripts and Schemas. Many researchers have pointed out that computationally fixed mental structures such as scripts are too inflexible to serve the purposes for which they were originally designed [e.g., Schank 1985, van Dijk & Kintsch 1983]. Barsalou [1987] has found that the variables active for a schema vary across contexts in which a schema is used. Similarly, the different tracks that people construct for scripts appear to vary widely across contexts. This may be because these structures exist only as temporary constructs in working memory rather than as invariant structures in long-term memory. Research by Kintsch and Mannes [1987] further supports the idea that scripts do not exist as invariant structures in long-term memory. They claim that knowledge is not pre-organized in terms of scripts and schemata, but that such structures are generated from an unorganized associative net in response to a specific task demand in a specific context. Only in this way, they believe, can the flexibility and context sensitivity that characterize human script use be achieved.

2.4. Summary. An active, flexible representation of generic concepts seems necessary to support the empirical evidence that (1) graded categories, generic concepts, and scripts are not invariant structures; i.e., they vary with context. Our representations and processing are based on an active model of generic concepts.

3. THE STRUCTURE OF BASIC LEVEL CONCEPTS.

3.1. Parts and Basic Level Categories: Physical Object Categories. Tversky and Hemenway [1984] have demonstrated that the basic level differs qualitatively from other levels in taxonomies of objects and living things: part terms predominate in subjects' listings of attributes characterizing objects at this level. Parts are rarely used to characterize superordinate level categories, and members of different subordinate categories share parts and differ on other attributes. Thus, it appears that our knowledge at this level is organized around part-whole divisions. Berlin has also suggested that the perception of overall part-whole configuration is the fundamental determinant of the basic level [1978].

3.2. Perceptual Grounding of Basic Level Concepts. Basic level concepts consist of more than mere attribute lists of parts that we can talk about, however: they are perceptually grounded concepts. Thus, it seems clear that a geometric (3D) model of a basic level object's shape is also encoded in long-term visual memory so that it may be recognized on subsequent occasions [Marr 1982]. Physical objects have their own intrinsic axes used in the 3D model to determine the internal layout of parts: e.g., an up/down axis, a front/back axis, the center/periphery, the interior/exterior. This geometric model enriches these concepts. Lakoff characterizes this type of geometric information as image schematic structure. In discourse, we frequently find references to this image schematic structure; e.g., to the *front* of a house (front/back image schema), the *interior* or *exterior* of a car (interior/exterior schema), and the *top* or *bottom* of many objects (up/down schema).

Additional percepts, e.g., sounds, colors, odors, also characterize objects at the basic level. Thus, if a *dog* is mentioned, its *bark* is implicitly evoked, just as part-whole structures such as *legs, ears,* a *nose*, and a *tail* are evoked.

We depend on both perceptual and functional qualities to differentiate basic level objects into classes. Perceptual properties that have high diagnosticity may be salient, since they are useful for distinguishing instances of a concept from instances of other concepts. Functional properties relevant to how people typically interact with instances of a concept are likely to be highly salient as well.

3.3. The Representation of Basic Level Concepts in SNePS. We use default generalizations [Peters & Shapiro, 1987a] to represent facts about the typical exemplars or members of a category. Thus, a basic level concept in our semantic network is, in part, a collection of default generalizations about typical exemplars: generalizations about their (1) part-whole structure and image schematic structure that is derived from a geometric model in visual memory, (2) other perceptual structure, and (3) functional attributes. These attributes and structures are all useful in categorization, i.e., in identifying category members. This knowledge forms the context-independent structure of basic level concepts that underlies the use of these concepts in language in any context.

In addition, however, basic level concepts are connected to thematic associates (concepts related by events rather than by taxonomic similarity) and to other non-centrally related concepts (e.g., to attributes not used or useful in categorization). Thus, they are understood with respect to our knowledge and theories about all other connected concepts. The thematic associates and other non-centrally related concepts form the context-dependent structure associated with a concept, i.e., structure that is relevant in particular situations. For example, *bottle* is a thematic associate of *baby*, and *mortgage* is non-centrally related to *house*. Thus, each concept in our system consists of large amounts of interrelated knowledge; however, we use no high-level knowledge organization structures such as frames or schemas. Instead, we use this relatively unorganized, interrelated knowledge to construct "concepts" appropriate to the current task and context in working memory during discourse processing, i.e., to construct the appropriate frame or schema.

We understand basic level concepts such as *dog* and *skunk* not only in terms of appearances and affordances, but also with respect to our deeper theories about *animals*. Thus, our understanding of basic and subordinate level concepts also depends upon our understanding of superordinate level concepts. In experimental studies, subjects frequently list few or no attributes for these concepts [Tversky & Hemenway 1984]. Thus, our knowledge about superordinate concepts is a deeper and less easily verbalized knowledge, involving underlying principles and theories about the world. The representation of superordinate level concepts is a current area of research for us. We believe that it is an extremely important area.

4. DISCOURSE PROCESSING: IMPLICIT FOCUSING.

The processing performed by our system will be illustrated by considering the issues and problems in comprehending references to implicitly evoked entities in discourse. The problem of implicit focusing is an important one for natural language understanding systems, since references to entities not previously explicitly mentioned occur frequently in discourse. As stated previously, basic level concepts evoke a rich set of entities which may be referred to later in the discourse. Superordinate level concepts, in contrast, evoke few discourse entities.

4.1. Focusing Mechanisms. We have extended the grammar of SNePS/CASSIE [Shapiro 1982, Shapiro & Rapaport 1987] to handle implicitly focused definite anaphors. Resolution of references to explicitly focused entities makes use of a primary focus space; i.e., referents or co-specifiers are sought among the active elements of this focus space. Implicitly focused entities, however, require a second focusing mechanism, which makes use of an additional data structure called a *potential focus list* and the SNePS path-based inference package [Shapiro 1978].

4.2. Potential Focus List. When a new individual identified by its basic level name (e.g. a *dog*) or a generic basic level concept (e.g., the type *dogs*) is encountered in input, the context-independent satellite entities implicitly evoked by the central concept are placed in a potential focus list with the evoking concept. I.e., we believe that these reflex, or subconscious, inferences are made at the time of reading/hearing the central basic level concept. The kinds of inferences that are made are those discussed above: inferences about part-whole and other image-schematic structure, and inferences about perceptual structure. (Inferences about functional attributes are not currently activated, since the current discourse understanding task is concerned only with definite anaphora. Functional attributes are clearly a central component of basic level concept structure, however.) Inferences concerning context-dependent entities, i.e., thematic associates and other non-centrally related associates, are also made at this time.

When a subordinate level concept, e.g., a *Volkswagen*, is encountered, the reflex inference that *Volkswagens are cars* is made, and the subconscious inferences about *cars* are again drawn. Evidence suggesting that subordinate level concepts automatically activate their superior, basic level concepts has been provided by many researchers including Rosch [1978] and Barsalou [1982]. Thus, many of the subconscious inferences made when two different subordinates of the same basic level category are encountered are the same. This is supported by the inability of subjects to cite different properties of the sort discussed here, e.g., part attributes, to distinguish these subordinate categories [Tversky & Hemenway 1984; Rosch 1976]. E.g., *collies*, *poodles*, and *spaniels* share parts and differ on other types of attributes. Named individuals, such as *Lucy* (member of basic level category, *girl*) or *Rover* (member of basic level category, *dog*) are handled in the same way; i.e., the basic level concepts for these individuals are activated and the subconscious inferences about *girls* or *dogs* are made.

4.3. Path-Based Inference. The SNePS path-based inference package provides the subconscious reasoning that is required for implicit focusing. If a relation does not exist explicitly between two nodes in the network, one may specify a path of arcs which is semantically equivalent to the relation being sought. The definition of appropriate paths in the semantic network enables the automatic retrieval of the relevant satellite concepts of the basic level concepts: part structures, other image schematic structures, and other percepts.

Thus, entities such as parts can be retrieved by defining a path of arcs from a node representing a basic level category, e.g., *dog*, to its parts: *nose*, *tail*, *head*, *ears*, *legs*, etc. In addition to the parts path, we can define paths of arcs to retrieve percepts (sounds, colors, smells) and the other image schematic structures. The ATN grammar then makes use of these defined paths to activate, i.e., implicitly focus, the context-independent satellite concepts of a basic level concept (that has been either encountered in input or activated by a subordinate level concept): returning all the nodes (concepts) that are found at the end of the defined paths of arcs emanating from the basic level concept and placing them in the potential focus list.

Since basic level categories are represented, in part, in SNePS using default generalizations or rules [Peters & Shapiro, 1987a], for every default rule in the network that says that members of the basic level category typically or presumably have part P, P is put into the potential focus list. Figure 1 shows the default rule that can be paraphrased as 'For all x if x is a dog then typically x has a tail'' or more simply as ''Typically, dogs have tails''. It also shows a defined path called *parts* which is used to activate the sub-part entities of basic level concepts. These activated entities ''fade out'' of the secondary focus list when the evoking concept fades out of the primary focus space.

In a like manner, the context-dependent information may be implicitly activated. These concepts are not automatically activated whenever a basic level or subordinate level concept is encountered in input, however. Rather, their activation is triggered in response to a specific context. Currently, our implementation is limited to the activation of context-dependent information in response to discourse comprehension of events, such as *buying/selling a car/house, renting a car/house, driving a car, walking/washing the dog.*

4.4. Timing of Activation. Empirical evidence that reflex inferences about context-independent satellite entities and context-dependent thematic associates are made at the time of comprehension of concepts is provided by numerous studies, including work by Barsalou & Ross [1986], and Walker & Yekovich [1986]. In addition, Walker and Yekovich have demonstrated that it is not the case that all of the information associated with a concept is activated during comprehension of the concept. Rather, only the centrally related concepts become part of the discourse model in working memory, and, thus, are available as antecedents; peripherally related concepts are not available, i.e., are not implicitly activated. In our terminology, centrally related concepts are the context-independent entities and additional entities evoked by the current context; peripherally related concepts are the associates of a generic concept in long-term memory that do not become activated in the current context.

4.5. General Focusing Mechanism for Definite Anaphora Comprehension. When a definite anaphor is encountered, both the primary focus space containing explicitly mentioned entities, and the potential focus list containing implicitly evoked entities may be searched for a co-specifier, and the discourse model updated appropriately. I.e., if the co-specifier is found in the potential focus list, it is now moved to the primary focus space with the appropriate level of activatedness; if the co-specifier is found in the primary focus space, its level of activatedness is simply updated. If a co-specifier is not found in either the primary focus space or the potential focus list, then knowledge-based processing, i.e., inferencing using the total knowledge base, is necessary.

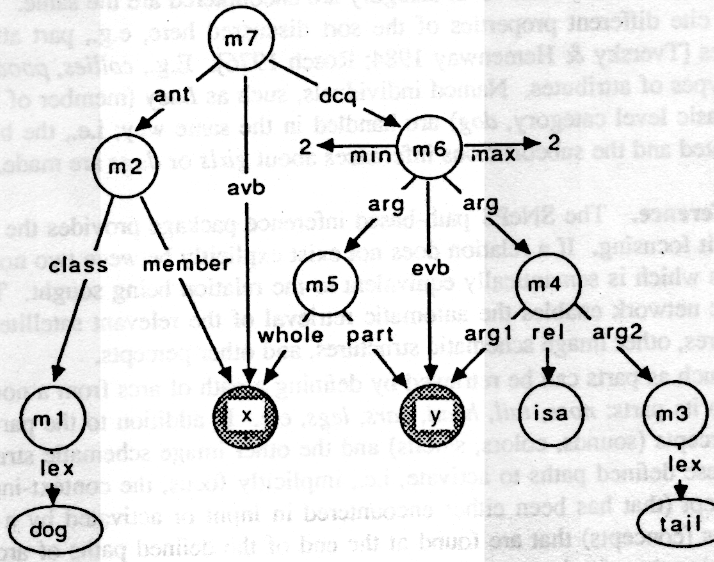

The following defines a path to find all the parts of basic level concepts

(def-path parts (compose arg2- arg1 part- whole avb- ant class))

Figure 1

(1) :Lucy bought a Victorian house.
I understand that Lucy bought a Victorian house
exec: 11.63 sec

(2) :The mortgage is high.
I understand that the mortgage of the Victorian house is high
exec: 8.90 sec

(3) :John visited her.
I understand that John visited Lucy
exec: 9.38 sec

(4) :The butler opened the door.
I understand that a butler opened the door of the Victorian house
exec: 17:45 sec

(5) :The house is huge.
I understand that the Victorian house is huge
exec: 7.30 sec

(6) :Lucy bought a dog.
I understand that Lucy bought a dog
exec: 9.35 sec

(7) :Her barking wakes Lucy.
I understand that the barking of the dog is waking Lucy
exec: 11.45 sec

(8) :John owns a cat.
I understand that John owns a cat
exec: 9.38 sec

(9) :He bought a canary.
I understand that John bought a canary
exec: 9.65 sec

(10) :The canary is named Tweety.
I understand that Tweety is the canary
exec: 6.40 sec

(11) :The cat stalks Tweety.
I understand that the cat is stalking Tweety
exec: 8.53 sec

(12) :His tail is swishing.
I understand that the tail of the cat is swishing
exec: 8.81 sec

(13) :His chirp alerted John.
I understand that the chirp of Tweety alerted John
exec: 11.61 sec

(14) :Lucy walked the dog.
I understand that Lucy walked the dog
exec: 11.56 sec

(15) :The leash became tangled.
I understand that the leash of the dog became tangled
exec: 10.86 sec

Figure 2

5. DEMONSTRATION.

Figure 2, a short dialogue from a sample run of SNePS/CASSIE, illustrates some of our current capabilities. User input is on lines with the :-prompt; the system's output and timing information are on the lines that follow. The following comments highlight some of the important features of this dialogue.

In sentence 1, comprehension of the basic level concept *house* implicitly evokes many entities, including parts such as *the roof, windows*, and *doors*. This forms part of the context-independent structure of the concept *house*. In addition, the verb *bought* in conjunction with *house*, activates such context-dependent entities as the *mortgage* and *cost* of the *house*. Thus, the concept *house* constructed in working memory has been tailored to the current context.

Sentence 2 contains a reference to an implicitly focused item, *the mortgage* which was activated in response to the context of *buying a house*. Sentence 4 contains a reference to *the butler*, an entity that was not activated by either the concept *house* or the current context. Thus, comprehension of *the butler* requires inferencing using the knowledge base. A comparison of the timing information for sentences 2 and 4 illustrates that the comprehension time for a non-activated entity *(the butler)* is longer than that for an activated entity *(the mortgage)*.

Sentence 6 contains the basic level concept *dog*, which implicitly activates many context-independent entities, such as parts and other percepts. It also contains the event of *buying a dog*, so *the cost* is also activated. In sentence 7, resolution of the definite anaphor *her barking* cannot be based on the normal mechanism of finding the most highly focused antecedent that matches the semantic features of the possessive pronoun *her*. It was not Lucy's barking! Rather, it requires a search of the potential focus list for the concept of *barking*, returning its evoking concept *dog*.

Sentences 12 and 13 also illustrate the need for using a focusing mechanism based on more than activatedness, recency and matching semantic features. I.e., *tail* was evoked as a part of the basic level concept *cat* (in sentence 7), but not as a part of the subordinate level concept *canary;* whereas *chirp* was evoked as associated with *bird/canary* (in sentence 8), but not with *cat*. Our processing of *his tail* and *his chirp* simply involves searching the potential focus list for these previously activated entities, not a search of the knowledge base. Since the potential focus list also contains the evoking concepts, integration of the activated associates (e.g., *chirp*) with the previously mentioned evoking concept (e.g., *bird/canary/Tweety*) is quite simple. Finally, in sentence 14, *walking the dog* implicitly activates many context-independent entities, as well as context-dependent associates such as a *leash*.

6. CONCLUSIONS.

We have presented representations for natural language processing that are based on (1) a recognition of the unique nature of basic level categories within natural category systems, and (2) the need for "active", flexible representations for generic concepts and categories. Thus, we have based our system on principles of categorization derived from current categorization research. We believe that most AI/NLP systems fail to take these principles seriously, and are, therefore, unable to support natural language processing in anything other than extremely limited domains.

REFERENCES

(1) Barsalou, L. W. (1982), "Context-independent and context-dependent information in concepts," *Memory & Cognition*, vol 10, pp. 82-93.

(2) Barsalou, L. W., & Ross, B. H. (1986), "The Roles of Automatic and Strategic Processing in Sensitivity to Superordinate and Property Frequency," *Journal of Experimental Psychology*, vol 12, pp. 116-134.

(3) Barsalou, L. W. (1987), "The Instability of Graded Structure," In U. Neisser (ed.), *Concepts and Conceptual Development* (Cambridge: Cambridge University Press).

(4) Berlin, B. (1978), "Ethnobiological Classification," *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 9-27.

(5) Bobrow, D. G., & Winograd, T. (1977a), "Experience with KRL-0, One Cycle of a Knowledge Representation Language," *IJCAI-77*, vol. 1, pp. 213-222.

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

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

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

(10) Kintsch, W., & Mannes, S. M. (1987), "Generating Scripts from Memory," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA (Hillsdale, NJ: Lawrence Erlbaum), pp. 861-872.

(11) Lakoff, G. (1987), *Women, Fire, and Dangerous Things* (Chicago: The University of Chicago Press).

(12) Marr, D. (1982), *Vision* (San Francisco: Freeman).

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

(14) Minsky, M. (1975), "A Framework for Representing Knowledge," In P. H. Winston, (ed.), *Psychology of Computer Vision* (McGraw-Hill).

(15) Minsky, M. (1986) *The Society of Mind* (New York: Simon & Schuster).

(16) Murphy, G. L., & Medin, D. L. (1985), "The Role of Theories in Conceptual Coherence," *Psychological Review*, vol. 92, pp. 289-316.

(17) Neisser, U. (1987), "From Direct Perception to Conceptual Structure," In U. Neisser, (ed.), *Concepts and Conceptual Development* (Cambridge: Cambridge University Press).

(18) Peters, S. L., & Shapiro, S. C. (1987a), "A Representation for Natural Category Systems," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA (Hillsdale, NJ: Lawrence Ehrlbaum), pp. 379-390.

(19) Peters, S. L., & Shapiro, S. C. (1987b), "A Representation for Natural Category Systems," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan (Los Altos, CA: Morgan Kaufman), pp. 140-146.

(20) 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.

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

(22) Schank, R. C., & Abelson, R. P. (1977), *Scripts, Plans, Goals, and Understanding* (Hillsdale, NJ: Lawrence Ehrlbaum Associates).

(23) Schank, R. C. (1985), *Dynamic Memory* (Cambridge: Cambridge University Press).

(24) Shapiro, S. C. (1978), "Path-Based and Node-Based Inference in Semantic Networks," In D. Waltz, (ed.), *Theoretical Issues in Natural Language Processing-2* (Urbana, Illinois), pp. 219-225.

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

(26) Shapiro, S. C. (1982), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *American J. Computational Linguistics*, vol. 8, pp. 12-25.

(27) Shapiro, Stuart C., & Rapaport, William J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network," in G. McCalla & N. Cercone (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer-Verlag), pp. 262-315.

(28) Tversky, B., & Hemenway, K. (1984), "Objects, Parts, and Categories," *Journal Of Experimental Psychology: General*, vol. 113, pp. 169-93.

(29) van Dijk, T. A., & Kintsch, W. (1983), *Strategies of Discourse Comprehension* (New York: Academic Press).

(30) Walker, C. H., & Yekovich, F. R., (1987), "Activation and Use of Script-Based Antecedents in Anaphoric Reference," *Journal of Memory and Language*, vol. 26, pp. 673-691.