

Superordinate and Basic Level Categories in Discourse: Memory and Context

Sandra L. Peters and William J. Rapaport
 Department of Computer Science
 State University of New York at Buffalo
 peters@cs.buffalo.edu, rapaport@cs.buffalo.edu


ABSTRACT

Representations for natural category systems and a retrieval-based framework are presented that provide the means for applying generic knowledge about the semantic relationships between entities in discourse and the relative salience of these entities imposed by the current context. An analysis of the use of basic and superordinate level categories in discourse is presented, and the use of our representations and processing in the task of discourse comprehension is demonstrated.

1. Introduction. We present representations for natural category systems based on a Roschian model of categories that has been extended to accommodate recent categorization research [Barsalou & Billman 1988; Keil 1989; Medin 1985, 1987; Murphy 1985, 1988, 1989]. We take issue with the assumption, implicit in most artificial intelligence (AI), natural language processing (NLP) systems, that generic concepts can be viewed as simple lists or collections of attributes. Richer representations of categories are needed to provide the intraconcept relations that structure categories and interconcept relations that provide connections to the rest of the knowledge base; these semantic relations provide some of the background or commonsense knowledge necessary for language interpretation. An analysis of the use of basic and superordinate level categories in discourse is presented, and the use of our representations and processing in the task of discourse comprehension is demonstrated. Published texts, the twenty English "Pear Story" oral narratives told by subjects after viewing a film [Chafe 1980], and forty unpublished narratives written by student subjects who were directed to retell the story of O. Henry's *A Retrieved Reformation* [unpublished data collected by Scott & Segal] provide the data analyzed in this paper.

2. Representations for Natural Category Systems. We have previously discussed the special status of the basic level in promoting inferences: the informativeness of the basic level arises from the large amount of information organized at this level and the perceptual grounding of basic level objects [Peters & Shapiro 1987ab; Peters, Shapiro, & Rapaport 1988]. In this paper we will discuss extensions to previously presented representations. Our implementation uses the SNePS knowledge representation and reasoning system, including a generalized ATN parser-generator [Shapiro 1978; Shapiro & Rapaport 1987]. Since the basic level has special status in our representations and processing, we will begin with this level.

3. Enhanced Representations for Basic Level Concepts. Default generalizations are used to represent facts about the typical members of a category in our system. Thus, a basic level category in our semantic network is, in part, a collection of default generalizations about part/whole structure, other image schematic structure, additional percepts, and functional and interactional properties [See Peters & Shapiro 1987ab, Peters, Shapiro, & Rapaport 1988 for a discussion of these structures]. Figure 1 shows a default rule that can be paraphrased as *for all x, if x is a car, then typically x has an engine*, or more simply as *typically cars have engines*. We build many such default generalizations about a basic level category such as *car*.



The following defines a path to find all the parts of basic level objects
 (def-path parts (compose arg2- arg1 part- whole forall- ant class))

Figure 1




Figure 2: Engines are interior parts of cars




Figure 3: Engines are mechanical parts of cars




Figure 4: Engines enable cars to run

Many researchers have pointed out that as people's knowledge increases, they come to reject mere collections of surface attributes and other typical features as being adequate to specify concepts; categories become further structured by "deeper" conceptual relations [Barsalou & Billman 1988; Keil 1987, 1989; Medin & Wattenmaker 1987; Murphy & Medin 1985]. Additional default rules are built to capture these "deeper conceptual relations. Thus, in addition to part-whole relations (*m5* in Figure 1) and relations about other percepts, we structure basic level categories such as *car* with enabling, functional, and spatial relations, such as those shown in Figures 2-4. (We have not shown the entire default rules, just the additional conceptual relations that provide intraconcept connections. I.e., *m8*, *m9*, and *m11* would replace *m5*, the part-whole relation, in the default rule of Figure 1, creating three similar default rules.) Figure 2 shows a spatial relation that further structures (and clusters) the interior parts of *car*. We structure the external parts of *car* similarly. Figure 3 is used to further structure or cluster mechanical parts of *cars*, such as the *brakes* and *engine*; Figure 4 shows an enabling relation: *engines enable cars to run/go*. Thus, in our system, there will be many assertions linking *car* and *engine*: the knowledge associated with a basic level category such as *car* is highly interconnected and organized by spatial, temporal, casual, explanatory, and enabling relations. Basic level categories in our system are highly structured by intraconcept relations. Interconcept relations provide additional structure. E.g., *mortgage* is a thematic associate of *house*. Figure 5 shows the enabling relation that connects these two concepts.

3.1. Context-independent and Context-dependent Structure. Our representations and processing are also based on the view that the information activated after hearing or reading a category name varies widely across linguistic contexts. I.e., categorizing an entity at the basic level provides access to a large amount of information; however, only a small subset of the information associated with a category in long term memory (LTM) is incorporated in the temporary concept constructed in working memory (WM) in a particular context. Barsalou [1982] has proposed that there are two kinds of information associated with categories in LTM: context-independent (CI) properties are activated by the word for a category on all occasions, independent of context; context-dependent (CD) properties are activated only in relevant contexts.

We had originally decided that topographic structure, i.e., parts that define the overall shape of basic level objects, were context-independent attributes, i.e., automatically activated across all contexts. [Peters, Shapiro, & Rapaport 1988]. However, an examination of the normative data showing the properties listed by subjects for basic level objects in free articulation tasks [Ashcraft 1978; Rosch et al. 1976; Tversky & Hemenway 1984] disconfirms the hypothesis that all of the external parts that contribute to the overall shape are context-independent. Production frequency of properties is considered to be a measure of semantic relatedness between category names and their properties [Ashcraft 1978], and although some exterior parts of objects are always generated in these tasks, many are not. In addition, some interior (hidden) parts and some non-part attributes are always articulated. E.g., subjects generate *wheels*, *tires*, *seats*, *engine*, *steers*, and *transportation* for *car*. Here, only *wheels* and *tires* are exterior parts. *Tail*, *long ears*, *white fur*, *soft*, and *animal* are generated for *rabbit* (only *tail* and *long ears* are external parts), and *wings*, *beak*, *feathers*, *flies*, *eggs*, and *nests* for *bird*. We hypothesize that context-independence arises as additional causal and explanatory relations integrate or interconnect these attributes and the category name that evokes them. Thus, *engine* achieves context-independence because of its functional importance and because of the many conceptual relations that interconnect *engine* and *car*. Many interactional properties are also CI; e.g., *seats* and *steering wheel* achieve context-independence because we interact with cars by sitting on the *seats* and using the *steering wheel*. Thus, not all parts may achieve CI status, but rather, one could argue that only parts attended to get processed in a manner to produce CI status; in this case, causally relevant parts would have a distinct advantage [Barsalou 1989, personal communication].




Figure 5: Mortgages enable buying houses




Figure 6: Having an engine is a context-independent property of cars

We currently create an assertion that marks CI properties, after a high degree of connectivity arises between properties and the category name. Figure 6 shows such an assertion, which we paraphrase as *having an engine is a CI property of cars*. CI properties, marked by such an assertion, are always activated when the category name is mentioned.

The category knowledge composed of the less strongly associated attributes (thematic associates and other non-centrally related entities) forms the context-dependent structure of basic level categories. We use the semantic relationships between verbs and nouns in discourse to evoke context-dependent entities, i.e., making the interrelated inferences normally made when particular verbs are used with particular nouns. E.g., *car* is associated with many different kinds of generalized actions: *driving, getting gas, washing, repairing, traveling, buying*. In our system, different knowledge associated with the category *car* is activated for each of these generalized actions. The demonstration runs in Appendix A show the activation of CI and CD entities in discourse comprehension, i.e., what interrelated inferences are made, and when they are made.

3.2. Examples of Contextual Fluctuations. Two short examples illustrate the different entities highlighted in different contexts. In the first passage, parts of the car relevant to driving, starting, and stopping are referenced:

As the *car* crept up the slope of the bridge, the inspector burst out laughing. He laughed so hard he could scarcely give his next direction. "Stop here," he said, wiping his eyes, "then start 'er up again."

Marian pulled up beside the curb. She put the car in neutral, pulled on *the emergency*, waited a moment, and then put the car into gear again. Her face was set. As she released *the brake*, her foot slipped off *the clutch pedal* and *the engine* stalled (A. Gibbs, *The Test*, p. 255, italics added).

In the second passage, in which the car is stopped, and the characters get out and move around the outside of the car, exterior parts of the car are referenced:

Lloyd brought the *car* to a screeching halt at the very edge of the boat launching ramp.

Lloyd looked calm, cool, and loaded as he got out and swayed his way around to *the rear* of the car. I thought he was going to open *the trunk* and present me with my father's coat, but instead he hoisted himself up on *the back* of the car. He lit a cigarette and from his perch on *the roof* he seemed to be enjoying my fright over the near plunge into the water (P. Zindel, *Confessions of a Teenage Baboon*, pp. 106-113, italics added).

The next example shows the use of activated CI information. Since *fur* is automatically activated for Cyril, the cat, (*fur* has CI status, since it is an interactional property; i.e., people pet cats, stroking their fur), disambiguation of the reference *his fur* is easily handled. Disambiguation here does not involve choosing the most activated or highly focused entity, but rather choosing an activated entity with this activated property. Claws also has CI status for cats.

When Jury {a detective} opened the door to Racer's {Jury's boss} sanctuary, Cyril {cat} slid between *his feet*, streaked snake-like across a carpet the color of *his fur*, and was scaling the bookcase set back against the wall to the left of Racer's large desk. *His claws* were like pinions digging into forensic science, ... (M. Grimes, *The Five Bells and Bladebone*, p. 188, italics added).

An example showing a reference to a CD entity follows. Here *the paw* is evoked in the context of the cat *striking* at the water:

The cat walked out of the secluded garden and toward the bank of a stream farther on. Here it crouched and watched a wren having a dust bath. Before it could pounce, the wren was away, skimming across the water. Looking into the stream, as if the bird might have fallen there, the cat saw shadows deep inside darting, hanging suspended, darting forward again. The cat struck at the water, trying to fix the moving shadow. ... It yawned again, washed at *the paw*, stopped when it saw something skittering across the footbridge and followed. (M. Grimes, p. 11, italics added)

4. Superordinates and Basic Level Names in Discourse. Since basic level categories carry the most information, they denote referents at their level of "usual" utility; i.e., at a level that is both sufficiently informative and cognitively efficient to manipulate. Thus the use of basic level terms to refer to entities in discourse follows the Gricean

maxims of conversational quantity and manner, since basic level names usually carry sufficient information for an addressee/reader to be able to identify the individual or category being referred to. What constraints cause a speaker/writer to abandon basic level terms in categorizing a given entity at a given point? In particular, when are superordinate level names used in discourse, and how is the knowledge stored with superordinate level concepts used during discourse comprehension? The next sections will analyze the use of superordinate terms in written and oral narratives.

4.1. Discourse Analysis: Use of Superordinate Level Names. Basic level terms constituted 93% of the nominal references to concrete entities in the twenty "Pear Story" oral narratives; superordinate level terms were used very infrequently, constituting only 2% of the references [Downing, 1980]. Basic level terms again predominated in the Scott & Segal subjects' written narratives, constituting 75% of nominal references to concrete entities. Superordinates occurred much more frequently than in the oral "Pear Story" narratives, constituting 17% of the references in the Scott & Segal written narratives. In the next sections superordinate level references will be examined more carefully.

4.1.1. Groups, Collections, and Classes. It has been suggested [e.g., Murphy & Wisniewski 1989; Wisniewski & Murphy 1989] that superordinates are frequently used to refer to groups, collections, and classes. I.e., basic level names are frequently used to refer to individuals or single objects, while speakers use superordinate names such as *furniture*, *clothes* and *plants* to refer to groups of related objects simultaneously [Wisniewski & Murphy 1989]. Our text analysis confirmed this usage of superordinate names. The following passages from O. Henry's *A Retrieved Reformation* (with italics added to indicate the superordinate level term) illustrate this usage:

(1) Take him back, Cronin, smiled the warden, and fix him up with outgoing *clothes*.

(2) Pulling out from the wall a folding-bed, Jimmy slid back a panel in the wall and dragged out a dust-covered suitcase. He opened this and gazed fondly at the finest set of burglar's *tools* in the East.

(3) He was at much at home in the *family* of Mr. Adams and that of Annabel's married sister as if he were already a member.

In the following sections we discuss additional uses of superordinate labels that were found in the written and oral narratives examined.

4.1.2. Introductory Mentions of Entities. Frequently superordinate labels are used to introduce entities to the hearer/reader, i.e., an author/speaker uses a superordinate level term for the first mention of an entity, later switching to a basic level term [Downing, 1980]. This technique is typically used only with characters and other elements that are central to the narrative. Two examples of this technique follow:

It was on the Dover road that lay, on a Friday night late in November, before the first of the *persons* with whom this history has business. [Dickens, *A Tale of Two Cities*, p. 8]

About three hundred eighty-five thousand years ago, when the oceans and continents were in place as we know them today, the land bridge from Asia was open, and a huge ponderous *animal*, looking much like an oversized elephant but with enormous protruding tusks, slowly made his way eastward, followed by four females and their young. [Michener, *Alaska*, p. 15]

In both of these examples the reader has the expectation that she will soon find out more about the characters being introduced by the superordinate level terms, and the superordinate level names seem to contribute to the readers being gradually "drawn" into the story. Thus, the readers' attention becomes focused on entities introduced in this way.

4.1.3. References to Focused Discourse Entities. A superordinate label is also frequently used to refer to a discourse entity that is focused, but not so highly activated that a pronoun or zero is easy to understand, yet too highly activated for a basic level term to be used without sounding childish or redundant. I.e., the basic level term may give too much information, given the focus level of the referent. For example:

He took a train for three hours, got off and went to a cafe owned by Mike Dolan. Dolan gave Jimmy his key. Jimmy went to a room in the back of the *establishment*. [Scott & Segal unpublished data, subject's retelling of O. Henry's, *A Retrieved Reformation*]

At the end of this passage the key and the cafe are competing, focused discourse entities (both can be referred to using *it*). Since the key is more highly focused, a pronominal reference to the cafe (use of *it*) seems less suitable than the use of the superordinate term *the establishment*. The use of the basic level term *the cafe* also seems less suitable than *the establishment*, because it is redundant. An additional example follows:

The cab stopped in front of Lloyd's house and I got out. I think the driver thought I was going to run away without paying but my slacks were so tight I couldn't get my money out of my pocket without standing up. I gave him the fare, including a good tip, which left me with about a dollar and twenty-nine cents to my name. It was the absolute end of

my savings from the last case my mother had, where her boss gave me a couple of dollars a day to walk his Yorkies. The driver just floored his junk heap and took off making so much noise I didn't hear the lilt of music until the taxi turned at the far corner. It had been quite a while since Helen and I had vacated the *place* and I couldn't believe the party was still going on. [P. Zindel, *Confessions of a Teenaged Baboon*, p. 141]

Here, a pronominal reference to the house (*it*) cannot be used, because there is too much distance between the first mention of the *house* and this reference. However, the superordinate, *the place*, is easily disambiguated. The next example again shows competing, focused discourse entities:

Anyway, this time it was a black cat reclining in the middle of this busy boulevard and it was positioned so it looked like it had been half run over by a car, squashing part of its body into a type of base like a buttress for a paper doll. It looked like I was staring at a stone statue of a cat with metallic eyes glazed with fire and that if I wanted to I could just take the *thing* home and use it as a garden ornament. The *animal* seemed dead except for its eyes, and I had the strangest feeling it might still be alive, . . . [P. Zindel, *Confessions of a Teenaged Baboon*, p. 121]

Here the superordinate label *the thing* refers to one discourse entity, the *stone statue of a cat*, while *the animal* refers to a second discourse entity, the *cat*.

4.1.4. Superordinates: Generality. Since superordinate labels are more general (less specific) than basic level terms, they can be used to communicate a sense of "vagueness". Thus, a superordinate name may be used to communicate that someone (a) doesn't know what something is, or what the name of something is (see example (1) below); (b) cannot see something very well (see example (2) below); (c) doesn't remember specifically what something was (see example (3) below); (d) doesn't want to call attention to irrelevant details or promote too many inferences (see example (3) again), or (e) wants to deliberately conceal information (e.g., in a mystery the detective may choose not to reveal the specific murder weapon used to commit the crime to a witness being questioned).

(1) Three boys came out, helped him pick himself up, pick up his bike, pick up the pears, one of them had a *toy*, which was like a clapper. And, I don't know what you call it except a paddle with a ball suspended on a string. [Chafe, 1980, transcript from English "Pear Film" narrative]

(2) Zoe squinted through a square of the lattice. Outside the crawl space it was lighter. She could see three *figures* standing beside a small truck. [Z. Oneal, *War Work*, p. 133]

(3) He [Jimmy] jumps on a train heading for the state boarder [sic] after he has some *food* and wine not of prison origin. [Scott & Segal unpublished data, subject's retelling of O. Henry's, *A Retrieved Reformation*]

In (1) the speaker telling the story of the "Pear Film" didn't know a name for the toy that a character in the film was playing with. In (2) Zoe cannot see the characters she is describing very well. In (3) the subject retelling the story either didn't remember that the food Jimmy Valentine had eaten was broiled chicken, or chose to be less specific when recounting this event, using a superordinate label rather than a more specific, basic level one.

4.2. Non-taxonomic Superordinates. In general, superordinates convey much less information than basic level terms. However, non-taxonomic superordinates (e.g., friend, enemy, jerk) frequently add information, communicating an attitude toward the referent(s):

I'm fifty and for half of those years I had worked for a Chicago newspaper I would rather not identify because it no longer exists. Its name exists, but it is a vulgar *rag* filled with sex and crime, edited by a gang of *creeps*, and owned by a right wing egomaniac who bought it for a song from the childish, idiotic woman who inherited it. [D. Kiker, *Murder on Clam Pond*, p. 22]

Non-taxonomic superordinates also frequently focus attention on particular attributes. In the following example, *enemy* evokes or activates entities such as *claws*, *tusks*, *teeth* as well as external body parts. All of these are relevant in the context of the struggle between these two enemies.

. . . he [Mastodon] heard a rustle that disturbed him. Prudently, he withdrew lest some enemy leap upon him from a hiding place high in the trees, and he was not a moment too soon, for as he turned away from the willow, he saw emerging from the protection of a nearby copse *his most fearsome enemy*.

It was a kind of tiger with powerful *claws* and a pair of frightful upper *teeth* almost three feet long and incredibly sharp. Mastodon knew that though this saber-tooth could not drive *those fearsome teeth* through *the heavy skin of his protected rear or sides*, it could if it obtained a secure foothold on *his back*, sink them into *the softer skin at the base of his neck*. . . Mastodon had *his long tusks*, of course, but he could not lunge forward and expect to impale his adversary on them, they were not intended for this purpose. [J. Michener, *Alaska*, p. 17, italics added]

4.3. Superordinates: Theories and Causal Reasoning. A number of researchers [e.g., Keil 1987, 1989; Murphy & Medin 1985; Barsalou & Billman 1988] have pointed out that our deep, rich theories about objects are linked to superordinate level concepts: theories that take us beyond categorization based on perceptual or surface similarity. In

order to appropriately categorize a whale as a mammal, we are required to move beyond using surface similarity as a basis for categorization: we need to recognize that hidden features and elaborated theories about origins are more important than perceptual features in categorizing biological kinds at a superordinate level. Theoretical knowledge also helps us to deduce reasons for correlations among basic level features: e.g., wings, hollow bones, small, and flies are correlated because they are causally linked. Our scientific knowledge organized at the superordinate level helps us understand these linkages. Taxonomic superordinates enrich basic level concepts, providing access to deep causal and explanatory relations.

4.4. Causal Reasoning in Discourse Understanding. Even when superordinates are not explicitly mentioned in discourse, they may provide the causal and explanatory relations necessary to our understanding of discourse. Our understanding of humor frequently depends on our use of superordinate level knowledge. E.g., in a short monologue about coffee, Garrison Keillor stated that although people said that decaffeinated coffee was dangerous and that a chemical used to decaffeinate coffee caused cancer in rats, his feeling was that the more rats that get cancer the better [Garrison Keillor's *American Radio Company of the Air*, February 16, 1990, WNED-FM]. This joke like many others is amusing because we recognize that an inappropriate conclusion has been drawn: the reasoning has been ineffectual. I.e., we know that rats, like people, are animals (in fact we know that both are mammals); that animal models are used to screen carcinogenic agents (with rats frequently being used as test animals); that if a chemical causes cancer in one kind of animal, there is a great likelihood that it will cause cancer in other animals (especially closely related animals). The knowledge used in this reasoning is superordinate level knowledge.

5. Discourse Processing: Inferences. The processing performed by our system will be illustrated by considering the problem of comprehending references to implicitly evoked CI and CD entities in discourse. As stated previously, basic level categories evoke a rich set of entities which may be referred to later in discourse. Superordinates, in contrast, evoke few discourse entities directly associated with the superordinate category itself, but frequently influence what is activated

The inferences made at the basic level seem to involve automatic processing: i.e., they are initiated by well-established memory structures, rely on connectivity between concepts, are computationally "cheap", and are plausible inferences rather than logically derived assertions which necessarily follow from the text. Superordinates, in contrast, do not seem to promote many inferences [Rosch 1981; Tversky & Hemenway 1984; Gelman 1989]: subjects frequently list few or no attributes for superordinate level concepts. Thus, our knowledge about superordinate concepts is a deeper and less easily verbalized knowledge, involving underlying principles and theories about the world. Inferencing at this level is not computationally "cheap" and effortless, but instead seems to require substantial attentional resources. The inferences made involve causal knowledge and explanatory relations (theoretically-oriented knowledge), and are deducible from the text and the knowledge base. We use a "spreading activation" form of inference similar to marker passing to activate the relevant basic level information, and deductive inference to reason about theory-laden superordinate level knowledge, forming assertions which follow logically or necessarily from the text.

5.1. Path-Based Inference. The SNePS path-based inference package provides the subconscious reasoning required for the automatic activation of CI and relevant CD entities: the definition of appropriate paths in the network enables the automatic retrieval of the relevant satellite concepts of basic level concepts. The context-independent entities can be retrieved by defining a path of arcs from a node representing a basic level category, e.g., *car*, to its context-independent properties: *wheels, tires, engine, seats, steering wheel*. The following path is used to retrieve all context-independent properties of basic level categories:

```
(find (compose arg2- arg- dcq- object- (domain-restrict (property context-independent)) object ant class)(find lex basic-level-category-name))
```

Our current implementation activates context-dependent information in response to discourse comprehension of events containing generalized actions such as *buying/selling a house, driving/stopping/repairing/washing a car, seeing/walking/washing a dog*. Additional paths are defined to retrieve these entities. Parts can be retrieved by defining a path of arcs from a node representing a basic level category, e.g., *car*, to its parts: *engine, roof, hood, trunk, wheels, brakes, etc*. Figure 1 in section 3 shows a defined path called *parts* which is used to activate these entities. The "parts" path is never used alone, rather it's used in conjunction with additional paths specified through default rules such as those shown in part in Figures 2-4, i.e., paths through the spatial, enabling, and other intraconcept and interconcept relations described earlier. E.g., we use an "exterior or surface parts path" to retrieve only exterior parts following generalized actions such as *washing a car* or *watching a bird*, and interior, mechanical parts following generalized actions such as *repairing a car*. Causal and enabling relations are considered to be more important than other relations in this system, and satellite associates that enable actions are activated by using "paths to enablers" to retrieve these associates in the relevant contexts. E.g., *mortgage* is retrieved following the generalized action, *buying a house*. An ATN grammar then makes use of the defined paths to activate, i.e., implicitly focus, the context-

independent and relevant context-dependent satellite entities of a basic level category (that has been either encountered in input or activated by a subordinate level concept): returning all the nodes that are found at the end of the defined paths of arcs emanating from the basic level categories and placing them in the system's working memory. (It is well established that categories automatically activate their superiors [Rosch 1978; Barsalou 1982].)

6. Demonstration of Discourse Processing. Appendix A shows four sample runs of SNcPS/CASSIE illustrating some of our current capabilities. User input is on lines with the **-prompt**; the systems' output and timing information are on the following lines. After a few of the sentences, a list of the activated context-independent (CI) and context-dependent (CD) entities or nodes (labeled CI evoked and CD evoked respectively) associated with the basic level category *car* is shown. (A complete listing of the associates of *car* activated in these sample runs is also found in Appendix A. See *Nodes Evoked in Demos 1, 2, & 3.*)

In sentence (1abc), comprehension of the basic level category *car* implicitly evokes many entities, including the following CI entities: the *engine*, *seats*, *steering wheel*, *wheels*, and *tires*.¹ In addition, the verb *bought* in conjunction with *car* activates a CD entity: the *price* of the *car*. Sentence (2a) of Demo 1 contains a reference to an implicitly focused CI item, *the engine*; sentence (3a), a reference to an implicitly focused CD item, *the price*. The verb *fixed* in conjunction with *car* in sentence (4a) of Demo 1 causes the activation of many CD associates of *car*, e.g., the *carburetor*, *distributor*, *spark plugs*, *battery*, and *brakes*.² Sentences (5a) - (7a) contain references to previously activated CI and CD entities.

In sentence (2b) of Demo 2, the generalized action *washing the car* causes activation of many CD entities: exterior parts of cars (the *bumper*, *grill*, *trunk*, *hood*, *roof*, *tires*, *wheels*, *windshield*, *finish*, and the *exterior*). Many of these activated entities are referenced in sentences (3b) - (7b).

In sentence (4c) of Demo 3, the mention of *starting the car* evokes many CD entities, e.g., the *ignition*, *key*, *battery*, *accelerator*, *brakes*, and a CI entity, the *engine*. Sentences (5c) - (12c) contain references to these activated entities.

In sentence (1d) of Demo 4 comprehension of the basic level category *house* implicitly activates many CI entities (e.g., *windows*, *doors*, *roof*, *shelter*). In addition, the verb *bought* in conjunction with *house*, activates such CD entities as the *mortgage* and *price*. Thus, the concept *house* constructed in WM has been tailored to the current context. Sentence (2d) contains a reference to the *mortgage* which was activated in (1d). Sentence (4d) contains a reference to *the butler*, an entity that has not been activated, since it is not a CI associate of *house*, and was not activated in the context of *buying a house*. Thus, comprehension of *the butler* requires inferencing using the total knowledge base. A comparison of the timing information for sentences (2d) and (4d) illustrates that in our system the comprehension time for a non-activated entity (*the butler*) is longer than that for an activated entity (*the mortgage*). Sentence (5d) contains a reference to the subordinate level category *collie*, which activates its immediate superior, the basic level category *dog*. The usual inferences about *dogs* are then drawn, and many CI entities (e.g., *tail*, *barking*, *animal*) are activated. It also contains the generalized action of *buying a dog*, so *cost* is activated. In sentence (6d), 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 WM for the concept of *barking*, returning its evoking concept *collie*. Sentences (12d) and (13d) 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 part of the basic level category *cat* in (7d), but not as part of either the subordinate level category *canary* or its immediate superior *bird*; whereas *chirp* was evoked as associated with *bird/canary* in (9d), but not with *cat*. Our processing of *his tail* and *his chirp* simply involves searching WM for these previously activated entities and their evoking concepts, not a search of the whole knowledge base. Thus, integration of implicitly evoked associates (e.g., *chirp*) with the previously mentioned evoking concept (e.g., *bird/canary/Tweety*) is quite simple and comprehension time quite fast.

¹ The CI entities that are always automatically activated in our system are those associates that have many causal and explanatory relations integrating them with the basic level category name that evokes them. They also are taken from property-norm data of Ashcraft, Rosch, and Tversky & Hemenway referenced earlier in this paper.

²Because of space limitations for these sample runs, only a few CD entities are shown as being evoked. Many more nodes are actually activated.

7. Future Research. Many problems remain to be solved. In particular, further work is needed in developing representations for and using superordinate level knowledge, in deciding which properties associated with a category have context-independent status, and in extending our work to consider context effects arising from the topic area, the task at hand, goals and discourse purposes.

REFERENCES

- (1) Ashcraft, M. (1978), "Property Norms for Typical and Atypical Items from 17 Categories: A Description and Discussion," *Memory & Cognition*, vol 6, no. 3, pp. 227-232.
- (2) Barsalou, L. W. (1982), "Context-independent and context-dependent information in concepts," *Memory & Cognition*, vol 10, pp. 82-93.
- (3) Barsalou, L. W. (1987), "The Instability of Graded Structure," In U. Neisser (ed.), *Concepts and Conceptual Development* (Cambridge: Cambridge University Press).
- (4) Barsalou, L. W., & Billman, D. (1988), "Systematicity and Semantic Ambiguity," In D. S. Gorfein (ed.), *Resolving Semantic Ambiguity* (New York: Springer-Verlag).
- (5) Downing, P. (1980), "Factors Influencing Lexical Choice in Narrative," in W. Chafe, (ed.), *The Pear Stories* (New Jersey: ALEX Publishing Company)
- (6) Keil, F. (1987), "Conceptual Development and Category Structure," In U. Neisser, (ed.), *Concepts and Conceptual Development* (Cambridge: Cambridge University Press).
- (7) Keil, F. (1989), *Concepts, Kinds, and Cognitive Development* (Cambridge: MIT Press)
- (8) Mervis, C. B., & Rosch, E. (1981), "Categorization of Natural Objects," *Ann. Rev. Psychol.*, vol. 32, pp. 89-115.
- (9) Medin, D., & Wattenmaker, W. (1987), "Category Cohesiveness, Theories, and Cognitive Archeology," In U. Neisser, (ed.), *Concepts and Conceptual Development* (Cambridge: Cambridge University Press).
- (10) Murphy, G., & Medin, D. (1985), "The Role of Theories in Conceptual Coherence," *Psychological Review*, vol. 92, pp. 289-316.
- (11) Murphy, G. L., & Wisniewski, E. J. (1989), "Categorizing Objects in Isolation and in Scenes: What a Superordinate is Good For," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 572-586.
- (12) Peters, S. L., & Shapiro, S. C. (1987a), "A Representation for Natural Category Systems I," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, WA, pp. 379-390,
- (13) Peters, S. L., & Shapiro, S. C. (1987b), "A Representation for Natural Category Systems II," *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, Milan, pp. 140-146.
- (14) Peters, S. L., & Shapiro, S. C. (1988), "Flexible Natural Language Processing and Roschian Category Theory," *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, Quebec, Canada, pp.
- (15) 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.
- (16) Rosch, E., & Lloyd, B. B., (eds.), (1978), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- (17) 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.
- (18) Shapiro, S. C., & Rapaport, W. 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.
- (19) Tversky, B., & Hemenway, K. (1984), "Objects, Parts, and Categories," *Journal Of Experimental Psychology: General*, vol. 113, pp. 169-93.
- (20) Wisniewski, E. J., & Murphy, G. L. (1989) "Superordinate and Basic Category Names in Discourse: A Textual Analysis," *Discourse Processes*, 12, 245-261.

CITED TEXTS

- (1) Dickens, C., *A Tale of Two Cities* (New York: Colonial Press)
- (2) Gibbs, A., (1940), "The Test," in *Short Stories from the New Yorker* (New York: Simon and Schuster), 252-256.
- (3) Grimes, M., (1987), *The Five Bells and Bladebone* (New York: Dell).
- (4) Kiker, D. (1986), *Murder on Clam Pond* (New York: Ballantine Books).
- (5) Michener, J., (1988), *Alaska* (New York: Random House).
- (6) North, S., (1963), *Rascal* (New York: E.P. Dutton & Co.).
- (7) Scott, P., & Segal, E., Unpublished Narratives, Retellings of O. Henry's, "A Retrieved Reformation".
- (8) Zindel, P., (1977), *Confessions of a Teenage Baboon* (New York: Harper and Row).

Appendix A

Demo 1: *Repairing the Car*

- (1a) : Lucy bought an old car.
I understand that Lucy bought an old car
CI evoked: (m138 m51 m20 m12 m3 m317)
CD evoked: (m94)
() exec: 12.41 sec gc: 4.71 sec
- (2a) : the engine was rebuilt.
I understand that the engine of the old car is rebuilt
() exec: 9.28 sec gc: 2.35 sec
- (3a) : the price was low.
I understand that the price of the old car is low
() exec: 9.21 sec gc: 2.33 sec
- (4a) : A mechanic fixed the car.
I understand that a mechanic fixed the old car
CD evoked: (m245 m224 m217 m153 m35 m3)
() exec: 13.08 sec gc: 2.43 sec
- (5a) : He tuned the engine.
I understand that the mechanic tuned the engine of the old car
() exec: 11.86 sec gc: 4.85 sec
- (6a) : He replaced the carburetor.
I understand that the mechanic replaced the carburetor of the old car
() exec: 11.81 sec gc: 2.38 sec
- (7a) : He fixed the distributor.
I understand that the mechanic fixed the distributor of the old car
() exec: 11.86 sec gc: 2.40 sec

Demo 2: *Washing the Car*

- (1b) : Lucy bought a red car
I understand that Lucy bought a red car
CI evoked: (m138 m51 m20 m12 m3 m317)
CD evoked: (m94)
() exec: 12.38 sec gc: 4.83 sec
- (2b) : She washed the car.
I understand that Lucy washed the red car
CD evoked: (m278 m265 m131 m87 m80 m73 m66 m59 m28 m20 m12)
() exec: 11.88 sec gc: 2.35 sec
- (3b) : She cleaned the windshield.
I understand that Lucy cleaned the windshield of the red car
() exec: 11.98 sec gc: 4.83 sec
- (4b) : She washed the grill.
I understand that Lucy washed the grill of the red car
() exec: 11.75 sec gc: 4.81 sec
- (5b) : She scrubbed the front bumper..
I understand that Lucy scrubbed the front bumper of the red car
() exec: 12.36 sec gc: 4.98 sec
- (6b) : She waxed the hood.
I understand that Lucy waxed the hood of the red car
() exec: 12.21 sec gc: 2.63 sec
- (7b) : The exterior is sparkling.
I understand that the exterior of the red car is sparkling
() exec: 10.55 sec gc: 2.48 sec

Demo 3: *Driving the Car*

- (1c) : Lucy bought an old car.
I understand that Lucy bought an old car
CI evoked: (m138 m51 m20 m12 m3 m317)
CD evoked: (m94)
() exec: 12.73 sec gc: 4.80 sec
- (2c) : she took a driving test.
I understand that Lucy took a driving test
() exec: 10.36 sec gc: 4.73 sec
- (3c) : She was nervous.
I understand that Lucy is nervous
() exec: 6.18 sec gc: 2.40 sec
- (4c) : she started the car.
I understand that nervous Lucy started the old car
CD evoked: (m224 m160 m147 m43 m35 m3)
() exec: 12.51 sec gc: 2.41 sec
- (5c) : she flooded the engine.
I understand that nervous Lucy flooded the engine of the old car ()
exec: 13.01 sec gc: 4.85 sec
- (6c) : she restarted the car.
I understand that nervous Lucy restarted the old car
() exec: 11.33 sec gc: 2.41 sec
- (7c) : she tapped the accelerator.
I understand that nervous Lucy tapped the accelerator of the old car ()
exec: 12.90 sec gc: 4.95 sec
- (8c) : she entered a busy street.
I understand that nervous Lucy entered a busy street
() exec: 12.30 sec gc: 4.86 sec

Demo 4: *General Demo*

- (1d) : Lucy bought a Victorian house.
I understand that Lucy bought a Victorian house
() exec: 11.28 sec gc: 4.83 sec
- (2d) : The mortgage is huge.
I understand that the mortgage of the Victorian house is huge
() exec: 9.28 sec gc: 2.35 sec
- (3d) : John visited her.
I understand that John visited Lucy
() exec: 8.48 sec gc: 2.48 sec
- (4d) : The butler opened the door.
I understand that a butler opened the door of the Victorian house ()
exec: 20.83 sec gc: 7.45 sec
- (5d) : Lucy bought a collie.
I understand that Lucy bought a collie
() exec: 11.83 sec gc: 2.61 sec
- (6d) : Her bark wakes Lucy.
I understand that the bark of the collie is waking Lucy
() exec: 13.98 sec gc: 5.16 sec
- (7d) : John owns a cat
I understand that John owns a cat
() exec: 10.20 sec gc: 2.46 sec

Demo 4 Continued

- (8d) : The cat is named Sylvester.
I understand that Sylvester is the cat
() exec: 8.05 sec gc: 0.00 sec
- (9d) : He bought a canary.
I understand that John bought a canary
() exec: 10.66 sec gc: 2.48 sec
- (10d) : The bird is named Tweety.
I understand that Tweety is the bird
() exec: 8.60 sec gc: 0.00 sec
- (11d) : The cat stalks Tweety.
I understand that Sylvester is stalking Tweety
() exec: 11.86 sec gc: 2.36 sec
- (12d) : His tail is swishing.
I understand that the tail of Sylvester is swishing
() exec: 10.63 sec gc: 2.53 sec
- (13d) : The chirp alerted John.
I understand that the chirp of Tweety alerted John
() exec: 13.46 sec gc: 2.58 sec

- (9c) : she approached an intersection.
I understand that nervous Lucy approached an intersection
() exec: 10.25 sec gc: 2.48 sec
- (10c) : she stopped the car.
I understand that nervous Lucy stopped the old car
() exec: 13.15 sec gc: 2.58 sec
- (11c) : the brakes were squeaky.
I understand that the brakes of the old car are squeaky
() exec: 8.25 sec gc: 0.00 sec
- (12c) : she failed the test.
I understand that nervous Lucy failed the driving test
() exec: 11.03 sec gc: 2.58 sec

Nodes evoked in Demos 1, 2, & 3:

(m3 (lex (engine)))
(m12 (lex (tire)))
(m20 (lex (wheel))) (m28 (lex (windshield)))
(m35 (lex (brake))) (m43 (lex (accelerator)))
(m51 (lex (seat))) (m59 (lex (trunk)))
(m66 (lex (roof))) (m73 (lex (hood)))
(m80 (lex (exterior)))
(m87 (lex (finish))) (m94 (lex (price)))
(m131 (lex (door)))
(m138 (lex (steeringwheel)))
(m147 (lex (ignition)))
(m153 (lex (carburetor)))
(m160 (lex (key)))
(m217 (lex (sparkplugs)))
(m224 (lex (battery)))
(m245 (lex (distributor)))
(m265 (lex (grill)))
(m278 (lex (bumper)))
(m317 (lex (vehicle)))