The Role of Correlated Properties in Accessing Conceptual Memory

Ken McRae
Virginia de Sa
University of Rochester, Rochester, NY

Mark S. Seidenberg
University of Southern California, Los Angeles, CA

keywords: correlated properties, conceptual memory, word meaning, connectionist models, semantic priming

Manuscript submitted to *Cognitive Psychology* in January, 1993.
Please do not quote without permission from the authors.

ABSTRACT

A fundamental question in research on conceptual structure concerns how information is represented in memory and used in tasks such as recognizing words. The present research focused on the role of correlations among semantic properties in conceptual memory. Norms were collected for 190 entities from 10 categories. Property intercorrelations influenced people's performance in both a property verification task and a short interval semantic priming experiment. Furthermore, correlated properties were more important for biological kinds than for artifacts. A connectionist model of the computation of word meaning was implemented in which property intercorrelations developed in the course of learning. The model was used to simulate the results of the two experiments. We then tested a novel prediction derived from the model: that the intercorrelational density of a concept's properties should influence the speed with which a concept is computed. This prediction was confirmed in a final experiment. We concluded that encoded knowledge of property co-occurrences plays a prominent role in the representation and computation of word meaning.

This work was supported by an FCAR (Quebec) doctoral fellowship and an NSERC (Canada) postdoctoral fellowship to the first author, an NSERC 1967 (Canada) doctoral fellowship to the second author, NIMH grant PO1-MH47566 to the third author, and NIH grant HD 27206 to M. K. Tanenhaus and G. Carlson. Part of this research formed KM's McGill University doctoral dissertation. We would also like to thank the McGill-IBM cooperative project in Science, Medicine, and Engineering for donating IBM microcomputers used in these experiments. We are also indebted to Mike Tanenhaus, Michael Spivey, and Kyunghee Koh for comments and helpful suggestions on earlier drafts.

Correspondence should be sent to Ken McRae, Department of Psychology, Meliora Hall, University of Rochester, Rochester, NY, 14627. email: kenm@psych.rochester.edu.

A central issue in cognitive science concerns how people store information about real-world entities, and how this knowledge is accessed when a person reads or hears a word. Many theories assume that concepts are represented in terms of properties (for an extended discussion, see Smith & Medin, 1981). For example, TIGER might be represented as *an animal* that *has fur*, *has claws*, *has whiskers*, and *lives in the jungle*.[1] However, a list of properties does not fully capture everything that people know about entities. For example, people also seem to encode knowledge of the relationships among an entity's properties (Malt & Smith, 1984). That is, people seem to encode that *has fur* and *has whiskers* are correlated because entities that have fur, such as tigers, dogs, and lions, also tend to have whiskers. A number of researchers (e.g., Boyd, 1984; Malt & Smith, 1984; Medin, 1989; Rosch, 1978; Rosch, Mervis, Gray, & Boyes-Braem, 1976) have claimed that an important step toward gaining a better understanding of conceptual memory is to determine how property intercorrelations affect the representation of concepts in memory.

Knowledge of the correlational structure of semantic space has figured prominently in discussions of the distinction between biological kinds (living things) and artifacts (objects that are made by humans). A number of studies have demonstrated basic differences in the ways that people process these two types of entities (e.g., R. Gelman, 1990; S. Gelman & Kremer, 1991; Keil, 1987). Boyd (1984) has argued that, to some extent, these differences arise because biological kinds are structured around sets of "contingently clustered" properties. Thus, he has claimed that the network of causally interconnected properties is much richer for biological kinds than for artifacts. Keil (1989) has argued that biological kinds have essences, where an essence is an underlying source of a dense network of property intercorrelations. In contrast, the properties of artifacts are much less causally connected because artifacts are structured around function; they do not have essences.

Recent computational modeling would also suggest that correlations among properties of entities may be important to conceptual representation. Recently, connectionist models have been successfully applied to simulating cognitive and neural processing (e.g., Brown, Kairiss, & Keenan, 1990; McClelland & Rumelhart, 1986; Rolls, 1990). These models feature correlational learning in that they naturally encode correlations among the properties of stimuli in their environment. This characteristic may play an important role in accessing concepts from words. Hinton and Shallice (1991) have noted that in many languages, including English, the mapping from word form (spelling or sound) to word meaning contains few regularities. Because of the dynamics of iterative connectionist models, some form of structure is required for computing stable patterns. Hinton and Shallice's model suggests that the correlational structure of semantic space provides this information. Thus, people's knowledge of how properties co-occur in entities may play a vital role in the dynamics of the computation of word meaning.

The research reported below used human empirical studies with real-world concepts to explore property-based conceptual representations, focusing on the role of property intercorrelations. These studies suggested a model in which encoded knowledge of property co-occurrences plays a key role in accessing conceptual representations from words. A connectionist model with this characteristic was implemented and used to simulate the human studies. A further prediction was derived from the model and verified in the final experiment.

Overview

This article is divided into eight sections. In the remainder of this section, theoretical assumptions underlying this research are outlined, and existing empirical investigations of the influence of property intercorrelations on categorization and recognition are briefly reviewed.

Section 2 contains three main subsections. First, a norming study is described in which subjects listed semantic properties for 190 entities from 10 categories. The second subsection outlines how these norms were used to construct conceptual representations in terms of individual

---

[1] When referring to entities in the world, lower case normal font is used. Concept names are printed in upper case, as are examples of experimental stimuli. Italics denote names of properties.

properties and correlated property pairs. These representations were used to unconfound the influences of individual and correlated properties when selecting stimuli and analyzing results. Experiment 2 demonstrates that correlations among properties affect subjects' performance when judging whether a property (e.g., *has an engine*) is true of an object (e.g., a motorcycle).

Section 3 deals with the role of property intercorrelations in concept similarity. In addition, differences between artifacts and biological kinds are explored, focusing on the relevance of property intercorrelations. We begin by describing how we computed similarity between two concepts from the conceptual representations of Section 2. In Experiment 3, similarity in terms of individual and correlated properties are used to predict the size of short interval semantic priming effects. It is shown that concept similarity as measured by individual properties predicts priming effects for artifact prime-target pairs, but concept similarity as measured by correlated property pairs predicts priming effects for biological kinds.

In the fourth section, a connectionist model is presented in which encoded information about property co-occurrence plays a key role in computing word meaning. Its architecture, training, and general performance are described.

In sections 5 and 6, we present simulations of the property verification and priming experiments, respectively. It is shown that the model provides a framework for plausible explanations of these results.

In section 7, a central characteristic underlying the model's performance is illustrated; a concept's convergence latency is partly determined by the intercorrelational density of its properties. This prediction is verified in Experiment 4.

The final section contains a summary of the main findings and a description of some limitations of the empirical and modeling work.

1.1 Background

The work described in this article focused on the representation and use of knowledge about simple real-world entities. This type of stimuli has dominated work on how real-world entities are categorized, presumably because it is possible to identify semantic properties relevant to their description (e.g., Rosch & Mervis, 1975; Smith, Osherson, Rips & Keane, 1988). The assumption that mental representations of entities are based on conceptual primitives or properties is incorporated by computational models in a number of domains, including: general memory models (e.g., Anderson, 1973; Hintzman & Ludlam, 1980; Murdock, 1987); models of conceptual representation (e.g., Anderson, 1991; Gluck & Bower, 1988; Hintzman, 1986; Kruschke, 1992; Medin & Schaffer, 1978; Neuman, 1974; Rosch & Mervis, 1975; Smith et al.); models of lexical representation (e.g., Dell, 1986; Hinton & Shallice, 1991; McClelland & Rumelhart, 1981; Seidenberg & McClelland, 1989); and models of visual representation (e.g., Beiderman, 1987). However, limitations of the property-based approach to categorization have been noted (e.g., Murphy & Medin, 1985; Palmer, 1978). Despite such limitations, the assumption that important aspects of our knowledge of the world can be represented by sets of conceptual primitives is integral to many models. The property-based approach remains central to the study of concepts and categorization because models within this framework have achieved considerable success in accounting for empirical facts, even if they are not relevant to the entire range of phenomena in these areas (e.g., Hayes-Roth & Hayes-Roth, 1977; Medin & Schaffer; Neuman; Rosch & Mervis; Smith et al.).

Probabilistic Representations

The experiments described below presuppose a probabilistic view of conceptual representation, where this view encompasses both prototype and exemplar theories. In most versions of prototype theory, a concept consists of the properties that are most typical or characteristic of the instances comprising the category (Medin, 1989; Rosch & Mervis, 1975; Smith et al., 1988). For example, the prototype concept BIRD consists of properties that are characteristic of birds, such as *has wings*, *has feathers*, and *flies*. Note that *flies* is not a necessary

or sufficient property of birds; penguins, which are birds, do not fly, but bats, which are mammals, do. Similar to prototype theory, exemplar theory is probabilistic (Hintzman, 1986; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1986); characteristic properties play a major role in representation and categorization. In exemplar models, each instance is stored in memory when it is attended; a category is represented by the instances that are labeled by the category name. For example, the concept BIRD would be represented by the set of birds that a person has experienced and stored in memory.

Recently, knowledge-based theories of categorization have been developed to account for complex categorization phenomena (e.g., Keil, 1989; Medin, 1989; Murphy & Medin, 1985). These theories typically assume a probabilistic component that is embedded in a rich and complex structure of pragmatic knowledge and internalized theories about the world. Although it has been shown that the content of the probabilistic component may be partly determined by people's intuitive theories about the world (Barrett, Abdi, Murphy, & Gallagher, 1992), the present article is not directly concerned with knowledge-based theories. We assume that the more that a task involves complex judgments or problem-solving, the more relevant is theory-based reasoning. In the experiments reported below, we emphasized rapid, on-line processes in which theory-based reasoning seems to be less relevant. In Experiment 2, for example, subjects were asked to decide if a property (e.g., *is crunchy*) was true of an entity (e.g., an apple). Experiment 3 concerns similarity between concepts within a level of the conceptual hierarchy (e.g., between DOG and CAT rather than between POODLE and DOG). Rather than using a slow, conscious similarity judgment task, a short interval semantic priming task was used in which the target (e.g., SKIRT) was presented 250 ms after the onset of the prime (e.g., DRESS). Given that short interval priming is considered to be automatic (Neely & Keefe, 1989), we assumed that the size of the priming effects resulted from property-based concept similarity. Finally, in Experiment 4, subjects decided as quickly as possible whether the thing to which a word referred is a concrete object in the world (e.g., APPLE versus MOTIVATION).

In summary, it was assumed that the experiments reported in this article tapped a probabilistic component of people's conceptual knowledge.

Defining Correlations Among Properties

There are three main ways that correlations among properties have been conceptualized. First, there are within-category instance-level correlations (Rosch, 1978). For example, a person's concept of DOG might consist of a set of properties (e.g., *has paws*, *has four legs*, *wags tail*, *barks*) because they tend to be strongly intercorrelated across the individual beagles and german shepherds that have been encountered. Storing this type of correlation may be regarded as encoding the properties that are correlated with the use of the category label, "dog" in this case. Notice that, given this perspective, a statistical relationship between two properties would be computed using information within a category.

As has been pointed out by Murphy and Wisniewski (1989), in most of the work reported on correlated properties in artificial concepts, what researchers have termed correlated properties have corresponded to conditional probabilities (e.g., the widely cited Experiment 4 of Medin, Altom, Edelson, & Freko, 1982). In these experiments, a subject is typically asked to learn to discriminate between two categories. A study might focus on two properties that occur in both categories. The properties might perfectly co-occur in instances of one category and never co-occur in instances of the other; therefore, across all instances, the properties occur independently. Rather than testing whether subjects are sensitive to property intercorrelations, these studies test whether people can use a pairwise cue to discriminate between categories.

In contrast, the work presented below focused on correlations among properties across real-world entities. Thus, two properties were considered as correlated if they tended to co-occur when all entities were taken into account, not just in entities of a specific category. Our primary concern was whether people encode correlations among properties in their environment, rather than if they can use property-pair information in a categorization or discrimination task. It should also be noted

that we measured correlated properties in terms of a statistical relationship, specifically, the Pearson product moment correlation.

1.2 Empirical Investigations of Property Intercorrelations

Until recently, most models of concepts and categorization tended to conceive of properties as independent and additive (Collins & Loftus, 1975; McCloskey & Glucksberg, 1978; Rosch & Mervis, 1975; Smith et al., 1974; Tversky, 1977, but see also Anderson, 1991). However, in recent years, within the domain of artificial concept research, evidence has accumulated suggesting that relationships among properties have important implications for categorization (Malt & Smith, 1984; Medin et al, 1982; Medin & Schwanenflugel, 1981; Murphy & Wisniewski, 1989, Experiment 3b; but also see Murphy & Wisniewski for failures to find effects of property intercorrelations), and for recognition of previously presented stimuli (Hayes-Roth and Hayes-Roth, 1977; Neuman, 1974; Younger & Cohen, 1983). In this section, two experiments are selected to illustrate influences of property intercorrelations. These two were chosen because they correspond to property intercorrelations as investigated in this article: correlational structure across the properties of all the instances that a person experiences (in the experimental situation).

Younger and Cohen (1983) studied ten-month old infants' ability to perceive correlations among properties in drawings of simple artificial animals. They constructed animals by manipulating five properties: *body type*, *tail type*, *shape of feet*, *shape of ears*, and *number of legs*. Each property could take on one of three values (e.g., *tail type*: *fluffy*, *feathered*, or *horse-like*). Infants were seated on their parents' lap and were visually habituated to animals constructed from combinations of property values. During habituation, *body type*, *tail type*, and *feet type* were perfectly correlated. Two of the three values on each of the five properties (e.g., *fluffy tail* and *feathered tail*) were presented with equal frequency during habituation; the third value was absent (e.g., *horse tail*). Test animals consisted of novel combinations of property values that either preserved or violated intercorrelations established during habituation. Other test animals consisted of property values not presented during habituation. Younger and Cohen designed two sets of habituation and test stimuli so that the correlated test stimulus for one group of infants served as the uncorrelated test stimulus for the other, and vice versa. They found that ten-month olds spent significantly less time viewing test animals that preserved property intercorrelations than those that violated them. Furthermore, viewing time for animals with completely novel property values was significantly greater than for test animals that preserved property intercorrelations, but similar to ones that did not. Younger and Cohen concluded that infants are sensitive to correlations among properties.

Medin et al. (1982) conducted four experiments to investigate people's sensitivity to correlated properties in a categorization task. They designed a simulated medical diagnosis task in which a disease was a set of five binary symptoms. In Experiment 3, in order to learn about a disease called burlosis, subjects studied case descriptions consisting of five binary, mutually exclusive symptoms (properties); for example, burlosis victims had either *puffy eyes* or *sunken eyes*. The cases were designed so that two symptoms were perfectly correlated and one value on each symptom was more characteristic (frequent). For testing, subjects were given pairs of cases and were asked to choose which was more likely to be suffering from burlosis. When both test cases violated the correlation, subjects tended to choose the one possessing a greater number of characteristic symptoms. In five test case pairs, a case containing a greater number of characteristic symptoms but violating the symptom correlation was contrasted with a case containing fewer characteristic symptoms but preserving the correlation. In these cases, cues to categorization based on property intercorrelations were directly opposed to those based on individual properties. Subjects were more likely to choose the test case with fewer characteristic symptoms and a preserved correlation. From these results, Medin et al. concluded that people's categorization decisions are influenced by individual properties and correlated property pairs.

In summary, studies such as these have demonstrated that children and adults encode correlations among properties of artificial concepts and use this knowledge in semantic tasks.

1.3 Property Intercorrelations in Real-world Entities

Artificial concept studies are important and interesting for a number of reasons: concept learning can be studied systematically, as well as variables such as frequency, recency, and similarity of exemplars; properties can be manipulated, as can relationships among category members. However, because of differences between artificial concept tasks and real-world situations, results from studies of artificial property intercorrelations may not apply to the real world. In an artificial concept experiment, a correlational structure across a set of exemplars is created by the experimenter. In the studies discussed above, properties were either independent or perfectly correlated. Because this manipulation is the strongest possible, it may have served to increase the salience of correlated property pairs. In the real world, where correlations among properties are presumably present at all strengths, distinctions like these would be rare, and correlations among properties may not be so obvious. Furthermore, we were interested in investigating the intercorrelations that an adult has encoded throughout a lifetime of experience with a huge number of properties. It is unclear how this experience corresponds to a limited exposure to a small number of instances that is typical of artificial concept learning experiments.

Curiously, although researchers have viewed property intercorrelations as a fundamental aspect of conceptual structure, only one study has addressed this topic using real-world concepts. Malt and Smith (1984) collected property norms for use in constructing conceptual representations for 15 entities from each of two artifact categories (CLOTHING and FURNITURE) and four biological kind categories (TREES, FLOWERS, FRUITS, and BIRDS). Each conceptual representation used all properties that were listed for two or more exemplars from within its category. A new set of subjects then rated applicability of each property to each entity (e.g., the applicability of *flies* to PENGUIN). Using the ratings, Malt and Smith computed the Pearson product moment correlation between each pair of properties within a category, finding that approximately 30% were significantly correlated, much greater than would be expected by chance.

In their main experiment, Malt and Smith (1984) tested for an influence of property intercorrelations on typicality judgments using three variants of Rosch and Mervis' (1975) family resemblance model. Family resemblance is defined in terms of the distribution of properties across exemplars within a category. Each concept received a family resemblance score that can be interpreted as representing its similarity to the category prototype or to other exemplars. This score corresponds to predicting typicality using independent and additive properties. Malt and Smith used property intercorrelations to predict typicality in two ways. For the additive pairs model, they assumed that correlated property pairs uniquely contribute to typicality judgments. To capture this assumption, for each significantly correlated pair, the two properties' family resemblance scores were added to a concept's score. Malt and Smith found that including additive pairs information in family resemblance slightly augmented predictions of typicality. The correlations between family resemblance scores and typicality with and without the correlated properties information, respectively, were: FRUIT (.81, .65); FLOWERS (.44, .23); TREES (.64, .33); BIRDS (.60, .63); CLOTHING (.56, .52); and FURNITURE (.79, .69). The largest increases in predictive ability occurred for three of the biological kind categories: FRUIT, FLOWERS, and TREES. Malt and Smith also used a subtractive pairs model to predict typicality. They assumed that, for a positively correlated property pair, if an exemplar had a high importance rating on one property, but a low rating on the other, then it violated the correlation (e.g., a PENGUIN *isa large bird*, but *has small wings*). If people are sensitive to correlational information, then violations may diminish an exemplar's typicality; however, they found no evidence to suggest this. In summary, although Malt and Smith found a high proportion of correlated properties in real-world entities, there was little evidence to suggest that property intercorrelations influence typicality judgments.

To summarize section 1, the correlational structure of semantic space has been shown to influence the learning and use of artificial concepts (e.g., Medin et al., 1982; Younger & Cohen, 1983). These demonstrations serve as an important constraint for models of concepts and categorization (e.g., Gluck, Bower, & Hee, 1989; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1984). However, results from artificial concept experiments may not generalize. The

sole reported investigation using real-world stimuli failed to uncover a measurable role for property intercorrelations (Malt & Smith, 1984). Therefore, the question of whether they influence how people learn and use real-world entity concepts remains unresolved.

2.0 EVIDENCE THAT PEOPLE ENCODE PROPERTY INTERCORRELATIONS

In section 2.1, Experiment 1 is described in which semantic property production norms were collected for 190 entities from 10 categories. Section 2.2 outlines how statistical analyses based on these norms were used to construct representations of concepts in terms of individual and correlated properties. Experiment 2 shows that people encode knowledge of the correlational structure of entities in their environment by demonstrating its influence on their performance in a property verification task. Stimuli were chosen for Experiment 2 using the conceptual representations constructed in section 2.2.

2.1 Experiment 1 - Property Norms

The method most often used to tap into people's knowledge of real-world entities is to have subjects list semantic properties (e.g., Hampton, 1979; McNamara & Sternberg, 1983; Rosch & Mervis, 1975; Smith et al., 1988). Property norms provide an empirical basis for determining people's probabilistic representations of entities. Their use is dictated by the fact that there are few other ways to gain systematic information about people's concepts. We did not assume that norms provide a literal record of conceptual representations. Rather, it was assumed that, because subjects use veridical conceptual representations to generate linguistically reportable properties for norms, the two are correlated. Thus, as Medin (1989) observed, properties that "are readily produced on feature listing tasks may not so much constitute the core of a concept as point toward it." (p. 1477)

In the first experiment, subjects were given a form containing 20 names of entities (e.g., DOG, DESK) and were asked to list properties for them (e.g., for DOG: *barks*, *has paws*). Space was provided for 10 properties. Subjects were asked to take a couple of minutes per object to list as many properties as possible.

Method

Subjects. Three hundred McGill university undergraduate and graduate students participated. In return for participation, a subject's name was entered into a lottery for a cash prize, or, in some cases, a subject was paid $2.

Materials. The stimuli were 10 category names plus 190 exemplar names. Nineteen exemplars were included from each of 10 categories. Exemplars of the 4 biological kind categories (BIRDS, MAMMALS, FRUITS, and VEGETABLES) and 6 artifact categories (CLOTHING, FURNITURE, KITCHEN ITEMS, TOOLS, VEHICLES, and WEAPONS) are listed in the Appendix. Categories and exemplars used by Rosch (1975) and Battig and Montague (1969) served as a guide for choosing the stimuli. To facilitate design of experiments reported below, items were chosen so that each category consisted of 9 pairs of similar exemplars (e.g., MOTORCYCLE-SCOOTER, SHIP-YACHT), plus a highly prototypical exemplar that could be paired with the category name (e.g., VEHICLE-CAR). The exemplars spanned a broad range of typicality.

Ten lists were constructed using the following criteria. Each list included 2 words from each of the 10 categories, 20 words in total. No list contained both members of a similar pair of concepts. Each list contained 1 category name. A subject received one list; therefore, 10 subjects were required to cover the complete set of 200 words. There were 300 subjects, 30 subjects per list. The pages within each list were constructed as follows. Each page contained 4 words. No page contained 2 words from the same category. The pages were randomly ordered. To minimize practice effects, a second set of lists was created using the following criteria. If 2 words appeared on the same page in the first set, they appeared on different pages in the second. If a word appeared in the top (bottom) row in the first set, it appeared in the bottom (top) row in the second.

Subjects were given as much time as needed to fill out forms. On average, it took approximately 40 minutes to complete a form, about 2 minutes per concept. Standard property norm instructions were included with each list of stimuli, followed by 3 example property listings.

Distribution and collection. Initially, 500 forms were distributed in psychology classes at McGill and 167 were returned to the laboratory, a return rate of 33%. In the second round, 200 forms were distributed and 80 were returned, a 40% return rate. The final 53 forms were collected in the laboratory; subjects were paid $2 each.

Recording the properties. Property descriptions from 30 subjects were collected for each exemplar and category name. For each item, a file was created that contained all semantically distinct properties listed for it. Recorded with each property was its production frequency, which is the number of subjects who listed the property for that particular concept (between 1 and 30). Thus, an item's description consisted of properties listed for it, with their associated production frequencies.

The initial stage in analyzing the norms was to ensure that synonymous properties were recorded identically, both within and between concepts. For example, *used for transportation*, *used for transport*, *is used for transportation*, *people use it for transportation*, and *transportation* were coded as *used for transportation*. It was equally important to ensure that properties differing in meaning had unique orthographic labels. The validity of all but the most obvious interpretations were verified with three naive colleagues. Subjects' responses were interpreted conservatively. If an interpretation was necessary, it was consistent across multiple instances of the response. The following method was adopted for interpreting and organizing the property set. Unless it was deemed important, quantifiers such as "generally", "usually", and "can be" were dropped. It was assumed that information carried by these quantifiers was inherent in production frequency data; that is, the number of subjects who included a property in their lists should vary according to how true a property is of an entity. If a subject listed an adjective-noun property (e.g., *has large ears*), it was divided (*has ears* and *has large ears*) under the assumption that the subject provided 2 bits of information: the entity in question has ears, and, those ears are large. Disjunctive properties (e.g., *is green or red*) were also divided (into *is green* and *is red*). A number of key words and phrases were used as leads for properties. These are displayed in italics, with an example completion in normal font: *a* deer (synonym), *bought/sold* in hardware stores, *causes* gas, *eaten* for dessert, *example* - Camaro, *found* in kitchens, *grows* underground, *has* paws (a part), *is* brown (adjective), *isa* tool (category), *lives* in the forest, *made* of metal, *part of a* table setting, *requires* a driver, *runs on* gasoline, *used* for carpentry, and *worn* by women. Finally, *has* was used as a key word for properties denoting a part (e.g., BULL - *has horns*), rather than for properties like BUS - *has a driver*, where a driver is not strictly a part of a bus in the way that wheels and seats are. The property *has a driver* for BUS was recorded as *requires a driver*.

Results

There were 57,458 responses, a mean of 9.6 properties per subject per concept. There were 19,739 distinct properties; on average, each property was listed by 3 subjects. 10,056 properties (approximately half) were listed for one concept only.

Subjects tended to perceive superordinate category names differently than exemplars. During debriefings, subjects commonly remarked that they found it "hard" or "strange" to list properties for "the words that don't really refer to anything", that is, the superordinate category names. In fact, subjects failed to provide coherent property descriptions for 8 of 10 superordinate categories; KITCHEN and BIRD were exceptions. These exceptions occurred presumably because KITCHEN is a place rather than a category name (KITCHEN ITEMS is actually the category), and BIRD is a basic level concept to most adults (Rosch et al., 1976). Subjects' inability to produce properties for superordinates can be illustrated by the percentage of responses that were exemplar listings. These percentages (in parentheses) were: MAMMAL (17), FRUIT (15), VEGETABLE (8), CLOTHING (14), FURNITURE (23), TOOL (20), VEHICLE (18), and WEAPON (23), which can be contrasted with BIRD (0.6) and KITCHEN (0.4). Because superordinates' property

listings were generally incoherent, they were excluded from all representations, analyses, and models reported below.

2.2 Property-based Representations

The norms were used to represent concepts in two ways: in terms of individual properties and in terms of correlated property pairs. These representations were constructed to explore implications of including information about property co-occurrence in conceptual memory. The individual properties representation can be conceived of as a representation of conceptual memory in which **no** information is encoded about co-occurrence of properties in the world, so that a concept is represented simply as a set of weighted properties. The correlated properties representation corresponds to a theory of conceptual memory in which people encode knowledge about how properties co-occur. This knowledge can be estimated through analyses of the production norms. Information about correlations among the properties of a concept is contained in a concept. The derivation of these representations is described in the next two sections.

Individual Properties Representation

An individual properties representation was constructed for the 190 entities from the property norms. Similar representations have been used by Malt and Smith (1984), Rosch and Mervis (1975), Smith et al. (1988), and Tversky (1977). All properties listed for any concept by a minimum of one-sixth of the subjects (5 or greater out of 30) were included. The resulting 1,242 properties were represented as vectors, so that concepts were represented as 1,242-unit patterns. Each property vector contained 190 units; the value for $unit_{ij}$ corresponded to the number of subjects who listed $property_j$ for $concept_i$. A sparse representation resulted; because no concept possessed more than 27 properties, each concept vector included a minimum of 1,215 zeros. In summary, in the individual properties representation, each concept was represented as a set of weighted properties.

Correlated Properties Representation

Concepts were also represented in terms of correlated property pairs. This representation resembled those used by Hayes-Roth and Hayes-Roth (1977) and Neuman (1974) because correlations among properties were explicitly represented. For example, in Neuman's property frequency model, when exemplars were experienced, a count was maintained of the frequency of individual properties and property pairs that they possessed. Similarly, Hayes-Roth and Hayes-Roth's model encoded the properties and property pairs that occurred in each exemplar, although all other elements of the power set were also encoded.

To construct the correlated properties representation, the Pearson product moment correlation was computed for each property pair. To avoid spurious correlations, only the 240 properties possessed by a minimum of 3 concepts were used. Thus, properties were correlated if they occurred together in the world as described by the norms. There was a vector in the representation for each of the 1,190 significantly correlated property pairs ($R^2 > .064$, $p < .01$).

A concept was represented as a pattern across 1,190 units. A value was assigned to each concept/property pair unit according to a three-part function. If a concept had neither property from the correlated pair, the unit was set to 0. If a concept possessed both properties, the unit was set to the sum of the production frequencies. If a concept contained only one property of a pair, then it violated the correlation and the unit's value was set to the negated production frequency of the property it possessed. Therefore, if a concept possessed only one property of a pair, the more often subjects listed that property for the concept, the more the correlation was violated. For example, the correlated properties *flies-has feathers* shared 42.5% of their variance. According to the norms, a CARROT neither *flies* nor *has feathers* (a value of 0); an EAGLE both *flies* and *has feathers* (13 + 16 = +29), and an OSTRICH *has feathers*, but does not *fly* (-22). In summary, in the correlated properties representation, a concept was represented as a pattern over all property pairs that were significantly correlated in the set of 190 normed entities.

In the analysis of property pairs, unlike Malt and Smith (1984), no significant negative correlations were obtained. This difference was due to the data on which correlations were computed. Malt and Smith used applicability ratings, so that if two properties consistently received ratings at the opposite ends of the scale within concepts, they were negatively correlated. In contrast, as mentioned above, the vectors of production frequencies that we used were sparse (a large proportion of units of a property vector were set to zero). Therefore, the presence of one property did not reliably predict the absence of another. However, as evidenced by the 1190 positively correlated pairs, the presence of one property did reliably predict the presence of another. This issue is discussed again in Section 4 when the connectionist model is described.

The individual and correlated properties representation are illustrated schematically in Figure 1. In the diagram, individual properties are illustrated as circles, and correlations between them as lines connecting the circles. Production frequency is not portrayed. Two concepts, labeled A and B, are illustrated by large ovals. A unit in the individual properties representation corresponds to a circle in Figure 1; a unit in the correlated properties representation corresponds to a connection between two circles. Therefore, a concept is represented either by a collection of properties, or a collection of connected properties. For example, concept A possesses 11 individual properties. It also possesses 9 correlated property pairs, and violates 3 others.

It should be noted that the individual and correlated properties representations were not connectionist models. Rather, they were conceptual representations that served in conjunction with statistical techniques to unconfound the influences of individual and correlated properties on people's performance on semantic tasks.

Figure 1 about here

2.3 Experiment 2 - Property Intercorrelations and Verification

We investigated whether intercorrelational strength among properties within a concept influences performance on a property verification task. In this task, a concept name, such as LION, was presented, followed by a target property name, such as *has teeth*. A subject was asked to indicate whether or not the target property was reasonably true of the object to which the concept name referred. For example, the correct response would be "yes" to LION - *has teeth*, but "no" to HORSE - *has scales*. A “reasonably true” rather than an "always true" criterion was necessary because, in line with the probabilistic view, few properties are "always true" of a concept (cf. Rosch & Mervis, 1975). For example, in an "always true" decision task, APPLE-*is red* and KNIFE-*is sharp* would require "no" responses.

Two groups of stimuli were created by pairing a target property with two concepts. For one concept, the target property was strongly intercorrelated with a number of other properties within that concept. In contrast, for the other concept, the target property was correlated as weakly as possible with other properties within it. The target properties in the two groups were equated for production frequency. We tested the prediction that if two groups are equated in terms of production frequency, then residual differences in verification latency can be attributed to property intercorrelations.

Method

Subjects. Twenty McGill university undergraduates were paid $2 each to participate.

Materials. The stimuli were 37 pairs of concepts, each sharing a target property (see the Appendix). For one concept from each pair, correlations between the target property and its other properties were as strong as possible (the "strong" group). For the second concept in each pair,

correlations between the target property and its other properties were as weak as possible (the "weak" group). For example, as shown in Figure 2, *hunted by people* was paired with DEER and DUCK. According to the norms, *hunted by people* is more strongly intercorrelated with DEER properties (11 of them) than with DUCK properties (4).

Figure 2 about here

Concepts were represented as in the individual properties representation. Rather than using the correlated properties representation directly, in this study, we computed the strength of intercorrelation for each target property-concept combination as the sum of the proportion of shared variances between the target property and all properties within a concept with which it was significantly correlated.

Thirty-seven target properties were paired with two concepts according to the following criteria. Intercorrelational strength was significantly greater for the strong group (mean = 175.4, standard error = 12.0) than for the weak group (mean = 19.7, standard error = 3.6), $t(36) = 14.42$, $p < .0001$ (one-tailed). No concept pair was used if the meaning of the target property might have been considered as different for the two concepts. For example, *has legs* has a different meaning for CANARY than for PANTS. Because a target property appeared in both groups, variables associated with it, such as reading time, were equated on an item basis. Also, production frequency was equated: strong (mean = 10.5, standard error = 4.5), weak (mean = 10.4, standard error = 4.9), $t(36) = 0.35$, $p > .4$ (two-tailed). Concepts were equated for familiarity, as judged on a 7-point scale by 20 subjects in a separate norming study: strong (mean = 5.36, standard error = 0.16), weak (mean = 5.36, standard error = 0.15), two-tailed $t(72) = 0.01$, $p > .9$. Finally, Ashcraft (1978) has shown that the number of properties produced for a concept in a norming task correlates with property verification latency; therefore, this factor was equated (strong: mean = 298, standard error = 4; weak: mean = 287, standard error = 5), two-tailed $t(72) = 1.70$, $p > .09$. Note that, although the difference might be considered marginally significant, it is less than 1/3 of a property per concept per subject.

In summary, on critical trials, a subject saw a target property preceded by a concept name. The properties in the strong and weak groups were equated in terms of production frequency, and the corresponding concepts were equated in terms of rated familiarity and number of properties listed in the norming task. However, in the strong group, the property was paired with a concept that contained a number of properties correlated with it, whereas in the weak group, the identical property was part of a concept that contained few correlated with it. This set of materials provided a means for testing whether intercorrelations among properties influence property verification.

Two lists were constructed so that no property was presented twice to a subject. Although 13 exemplars appeared with 2 target properties, repetitions were assigned to different lists so that subjects saw each exemplar only once. However, because PARAKEET appeared with 3 properties (twice in the strong group, once in weak), it appeared twice in list 1. Also, CARPET appeared with 4 properties (strong twice, weak twice) so that it appeared twice in each list. List 1 contained 18 items from the strongly intercorrelated group, and 19 items from the weak group. List 2 contained 19 items from the strong group, and 18 items from the weak group. List 1 contained 21 artifacts and 16 biological kinds. List 2 contained 22 artifacts and 15 biological kinds. To offset the 37 positive test items in each list, 37 negative filler items were prepared for which the correct response was "no, the property is not reasonably true of the concept". Type of property (e.g., part, function, or characteristic action) was approximately matched between test properties and negative fillers to prevent subjects from using it as a cue for responding. A further

10 fillers (5 positive, 5 negative) served as lead-in practice items. A unique set of 20 items (10 positive, 10 negative) were used in a separate practice session.

Procedure. In Experiments 2, 3, and 4, presentation of stimuli and recording of subjects' responses were accomplished using Micro Experimental Laboratory (MEL) software running on an IBM XT-286 microcomputer. A practice session with feedback was provided to ensure that subjects understood the "reasonably true" response criterion. Subjects were given an example in which they were told that the property *is curly* is reasonably true of HAIR, even though hair is not always curly. In contrast, *made of metal* is not reasonably true of HAIR. Subjects were instructed to press the "yes" key (always beneath their dominant hand) as quickly and accurately as possible if the property was reasonably true of the concept, or press the "no" key if it was not. There were 84 trials in the main section of the experiment. Ten lead-in items were presented in random order, followed by 74 test items, also presented in random order. Each trial proceeded as follows: a 1500 ms intertrial interval; an asterisk in the center of the screen for 500 ms; a blank screen for 100 ms; a concept name for 400 ms; the target property until the subject responded. It took approximately 5 minutes for a subject to complete the experiment.

Design. There was one factor, strength of the intercorrelations between the target property and other properties of a concept, with 2 levels, strong and weak. In the analyses, strength was a within subjects and a within items factor.

Results

Latency and accuracy of subjects' responses were recorded. One-tailed paired samples t-tests were used to determine the effect of intercorrelational strength on decision latency. All response latencies greater than 3 standard deviations from the mean of the correct test trials were replaced by the cutoff value (1.5% of the trials). Subjects were faster to judge that a property was part of a concept if there were strong intercorrelations with other properties in the concept (mean = 820 ms, standard error = 38 ms) than if there were weak or no intercorrelations (mean = 912 ms, standard error = 46 ms), $t(19) = 4.22$, $p < .0003$ by subjects, $t(36) = 4.23$, $p < .0002$ by items.

In linear regression analyses, intercorrelational strength was a unique and strong predictor of property verification latency. The dependent variable in the analyses was verification latency for the 74 items; the independent variables were intercorrelational strength and production frequency of the target properties[2]. Intercorrelational strength was a better predictor of verification latency, $R^2 = .18$, $F(1, 72) = 15.815$, $p < .001$ than was production frequency , $R^2 = .08$, $F(1, 72) = 6.408$, $p < .02$. Furthermore, both variables accounted for unique proportions of variance: intercorrelational strength $R^2 = .17$, $F(1, 71) = 14.805$, $p < .001$; production frequency, $R^2 = .07$, $F(1, 71) = 5.613$, $p < .03$.

Subjects committed errors on 8.6% of the test trials, and 4.2% of the negative filler trials. One-tailed paired samples t-tests were conducted using the square root of the number of errors as the dependent variable (Myers, 1979). Within the test trials, more errors were committed in the weak condition (13.0%) than in the strong condition (4.3%), $t(19) = 4.34$, $p < .0003$ by subjects, $t(36) = 3.24$, $p < .002$ by items.

Regressions were performed using square root of the number of errors as the dependent variable. Again, intercorrelational strength was a better predictor of verification performance, $R^2 = .12$, $F(1, 72) = 9.822$, $p < .003$ than was production frequency, $R^2 = .03$, $F(1, 72) = 12.273$, $p > .1$. Intercorrelational strength also predicted unique variance, $R^2 = .11$, $F(1, 71) = 9.034$, $p < .01$.

---

[2]Performing regressions in this manner essentially treats subjects as a fixed effect (Clark, 1973). The regressions do, nonetheless, represent a strong test because the norms on which the independent variables were based were collected independently of the priming study. The norms were conducted using a separate sample of subjects, and, did not, in any way, mention similarity among concepts. Recall also that no norming subject saw both members of a similar pair of concepts.

Discussion

The strength with which a property was intercorrelated with other properties within a concept influenced property verification. Thus, the results indicate that people encode knowledge of how properties of entities co-occur in the world. Mechanistic explanations of the source of these results are presented in Section 5, where simulations of this study are described.

3.0 CONCEPT SIMILARITY IN ARTIFACTS AND BIOLOGICAL KINDS

We now turn to the question of whether individual and correlated properties provide a basis for understanding similarity between entities. To test for an influence of property intercorrelations on concept similarity, a short interval semantic priming task was used. In this task, a prime was presented (e.g., LAMP), immediately followed by a target (e.g., CHANDELIER). The subject was instructed to read the prime, and respond overtly to the target. A priming effect is defined as the reduction in response time to a target when preceded by a similar prime (LAMP-CHANDELIER) versus a dissimilar prime (GOOSE-CHANDELIER). In Experiment 3, subjects made a semantic decision to each target (e.g., does the word refer to something that is made by humans?). It was hypothesized that the size of the priming effect for specific items could be predicted from concept similarity in terms of individual properties and correlated property pairs.

Concept similarity was measured by the cosine of the angle between two concept vectors. Cosine is a linear measure of similarity that ranges from -1 (opposite concepts) through 0 (orthogonal concepts, share no properties) to +1 (identical concepts). In the individual properties representation, each of the 190 concepts was a vector of production frequencies across 1242 properties (the number of properties that were listed for at least one concept by a minimum of five subjects). The cosine similarity measure increased linearly with number of shared properties, and decreased with number of distinct properties. According to the individual properties representation, SOFA and COUCH are very similar (cos = .879), PINEAPPLE and COCONUT are less similar (cos = .415), SCREWDRIVER and DRILL are marginally similar (cos = .214), and CHAIR and WHALE are dissimilar (cos = 0). Cosine between a concept pair expressed over individual properties was always positive because all values in this representation were positive (production frequencies).

In the correlated properties representation, a concept was represented as a pattern across 1190 correlated property pairs (in the statistical analyses of the properties, 1190 pairs were significantly correlated). According to this representation, SOFA and COUCH were quite similar (cos = .963), PINEAPPLE and COCONUT were slightly similar (cos = .219), SCREWDRIVER and DRILL were quite dissimilar (cos = .061), and CHAIR and WHALE were entirely distinct (cos = 0). The cosine between a concept pair expressed over correlated property pairs could be negative because violations were encoded.

A return to Figure 1 can provide a flavor of the concept similarity measure. In Figure 1, concepts A and B share 8 individual properties; three are distinct to A, and two to B. Roughly, in terms of individual properties, the concepts are more similar than they are different. However, concepts A and B share three correlated property pairs; two pairs are distinct to A, and one is distinct to B. Furthermore, A violates three correlated pairs of B, and B violates four pairs of A. Therefore, concepts A and B are much less similar in terms of correlated property pairs than in terms of individual properties.

The issue of correlated properties also bears on important questions regarding the relationship between artifacts and biological kinds. According to Keil (1989) and Boyd (1984), artifacts and biological kinds differ in the extent to which they entail correlated properties. They claim that biological kinds cohere around clusters of intercorrelated properties, but artifacts cohere around the intended function of the creator. We tested this idea by analyzing the individual and correlated properties representations. Two results suggested that correlated properties may be more important for biological kinds than for artifacts: (1) the properties of biological kinds are more densely intercorrelated, and (2) biological kind concepts possess far more correlated property pairs. By looking separately at biological kinds and artifacts, it was found that 11% of biological kind

property pairs were significantly correlated (i.e., correlated at $p < .01$), but only 6% of artifact pairs. More revealing, however, is the density of representations over those property pairs. The number of properties or correlated property pairs was used as the dependent variable in a two-way analysis of variance. The independent variables were type of entity (biological kind and artifact) and type of representation (individual properties and correlated properties). Using the 114 artifacts and 76 biological kinds that were normed in Experiment 1, it was found that type of entity interacted with type of representation, $F(1, 188) = 72.996$, $MS_e = 100.012$, $p < .0001$. Simple main effects indicated that biological kinds possess far more correlated property pairs (mean = 33.4, se = 2.1) than do artifacts (mean = 13.9, se = 1.2), $F(1, 367) = 145.884$, $MS_e = 118.483$, $p < .0001$. However, no difference was found between the number of individual properties possessed by biological kinds (mean = 16.5, se = 0.4) and artifacts (mean = 15.0, standard error = 0.3), $F < 1$. By category, the mean number of individual properties and correlated property pairs were: birds (15, 33), mammals (19, 38), fruits (16, 41), vegetables (15, 21), clothing (16, 19), furniture (15, 8), kitchen items (15, 12), tools (14, 7), vehicles (15, 14), and weapons (15, 23). Given that biological kinds possess more correlated property pairs than do artifacts, it may be that correlations are more important for biological kind similarity. If this is true, correlated properties may play a more prominent role in predicting priming effects for biological kinds than for artifacts.

Concept similarity as measured by overlap of individual properties reflects the linear aspects of similarity; similarity is a linear function of independent cues. In contrast, correlated property pairs introduce nonlinearities into the similarity measure in two ways. First, similarity over correlated properties results in a positively accelerating function over shared properties (see Shepard, 1957, 1986, for a discussion of the nonlinear, positively accelerating similarity/generalization gradient). To account for influences of multi-property cues on categorization and recognition, a number of models have incorporated mechanisms to make exemplar similarity a positively accelerated function (Gluck, 1991; Gluck & Bower, 1988; Gluck et al., 1989; Hayes-Roth & Hayes-Roth, 1977; Hintzman, 1986; Medin & Schaffer, 1978; Neuman, 1974; Nosofsky, 1986). The model of Neuman (1974) can be used to illustrate. In this model, property intercorrelations play a role in determining concept similarity because each additional shared property does not simply add a constant to the similarity measure. Given that the model encodes both properties and property pairs, if a pair of concepts share three properties, they have six units in common; if they share four properties, they have ten units in common; five units, fifteen in common. The effect of each additional shared property increases as similarity increases, resulting in a positively accelerating similarity gradient.

Concept similarity over correlated pairs introduces a second nonlinearity if violations of correlations are encoded. Over individual properties, similarity increases with shared properties and decreases with distinct ones. Likewise, over correlated properties, similarity increases with shared pairs and decreases with distinct pairs. However, similarity over correlated properties may also decrease when concepts share an additional property; similarity decreases when concept1 has both properties of a correlated pair (positive value) but concept2 has only one (negative value). Thus, when concept1 conforms to a correlation but concept2 violates it, they are less similar than if concept2 possessed neither property from the pair.

In summary, concept similarity in terms of correlated properties introduces nonlinearities into the measure. Given that property intercorrelations are dense for biological kinds but not for artifacts, similarity in terms of independent cues may be sufficient to measure artifact concept similarity, but the nonlinear aspects of the correlated properties representation may be required to estimate biological kind similarity.

3.1 Using Priming to Study Concept Similarity

Introducing this study is somewhat complicated by the fact that it is not standard procedure to use short interval semantic priming to study property-based similarity. This section is a slight digression in which we provide arguments concerning the advantages of this approach. Short interval semantic priming with a semantic decision task is an appropriate method to study concept

similarity for four reasons: it is influenced by memory internal processes (de Groot, 1984; den Heyer, Briand, & Dannenbring, 1983; Neely, 1977; Posner & Snyder, 1975); the influence of similarity is measured in an indirect manner as subjects do not intentionally judge similarity; and analyses are sensitive because a target can be used as its own control. Furthermore, in Experiment 3, semantic decision tasks such as "is it made by humans?" were used, ensuring that decisions must be based on semantic knowledge. These points are expanded below.

A number of studies strongly suggest that short interval priming is sensitive to effects associated with the computation of the meaning of the prime and target (de Groot, 1984; den Heyer et al., 1983; Neely, 1977; Posner & Snyder, 1975). Posner and Snyder proposed that, when a short time interval separates prime and target, semantic priming is automatic and memory internal. In contrast, with more time, subjects use the prime to generate an expectancy set for the impending target. This hypothesis has since been elaborated, but the core ideas have been retained (e.g., Becker, 1980; Neely & Keefe, 1989). Because the aim of Experiment 3 was to determine whether semantic priming effects could be explained in terms of property-based concept similarity, it was critical that priming effects were limited to the computation of concepts. A stimulus onset asynchrony of 250 ms was used because a number of studies have shown that this time interval is too short for expectancy-based priming to occur (Neely; den Heyer et al.; de Groot).

Second, it was assumed that an automatic process that is not based on a conscious comparison of two concepts is the best way to obtain a measure that is sensitive to property intercorrelations. Because they are slow and conscious, typicality and similarity judgments may be more dependent on individual properties or more complex types of knowledge that are central to knowledge-based theories of categorization (cf. Medin, 1989), minimizing effects of raw property intercorrelations. Existing empirical evidence supports this hypothesis. As described above, Malt and Smith (1984) found little evidence to suggest that property intercorrelations influenced typicality judgments. McRae (1991) found similar results using an explicit similarity judgment task (i.e., "How similar, on a scale of 1 to 7, are the things that these words refer to?"). The probability of uncovering influences of property intercorrelations may increase in a task where the influence of similarity lies in successive computation of two concepts.

Third, a priming paradigm allows for sensitive analyses because the critical comparison is between responses to a target word in two contexts. As detailed below, concept similarity was used in regression analyses to predict priming effects. Because a target was used as its own control, variance due to the decision component was factored out in these analyses. Furthermore, variance due to nuisance factors, such as target word length and frequency, was also removed.

In the vast majority of priming studies, subjects make lexical decisions to the target ("Is the target a word?"). Semantic decision tasks (e.g., "Is it an object?"), rather than lexical decision, were used to ensure that the decision was based on word meaning, and to allow a set of fillers to be constructed that removed the influence of retrospective priming. In contrast to semantic decisions, the extent to which lexical decisions are based on semantic information is not clear and is currently a matter of debate (Balota & Chumbley, 1984; Seidenberg & McClelland, 1989). Furthermore, in lexical decision, prime-target relatedness (similarity) can be used as a cue to the word/nonword response because the response to the target is "word" on 100% of the trials on which the prime and target are related. Neely and Keefe (1989) found evidence suggesting that priming effects in a lexical decision task are influenced by subjects retrospectively checking prime-target relatedness. In contrast, a retrospective check can be avoided in a semantic decision task such as "Is it animate?", because similar, but inanimate, prime-target pairs (e.g., HOUSE-COTTAGE) can be included as filler items. In summary, a semantic decision task was used because it must be based on computed word meaning, and, with appropriate fillers, similarity does not cue the decision.

In Experiment 3, broad semantic decision tasks were used because subjects were not likely to generate specific exemplars of a category that was as broad as all living things or all things made by humans. However, using a narrow semantic decision or categorization task (e.g., "is it a bird")

increases the probability that subjects can anticipate specific exemplars, possibly altering the basis of obtained effects (Jared & Seidenberg, 1992).

In summary, short interval semantic priming with a broad semantic decision task provides an appropriate and sensitive method for studying concept similarity.

3.2 Experiment 3

It was hypothesized that the magnitude of priming effects might be predicted using concept similarity in terms of individual and correlated properties. Because correlated properties introduce nonlinearities into the similarity measure and biological kinds possess densely intercorrelated properties but artifacts do not, it was further hypothesized that biological kind similarity might be captured by correlated property pairs, whereas artifact similarity might be captured by individual properties.

Method

Subjects. Forty-eight McGill university undergraduates were paid $3 each to participate.

Materials. The stimuli were 90 pairs of similar exemplars plus 10 category-exemplar pairs (see the Appendix). Note that Experiment 3 was conducted simultaneously with the collection of the property norms. At that time, it was not known that the norms would fail to provide reasonable category concepts. Therefore, the category-exemplar pairs were included in analyses of variance described below, but were excluded from regression analyses. There were 4 semantic decision tasks: "is it animate?", "does it grow?", "is it made by humans?", and "is it an object?". Similar and dissimilar prime-target pairs for each task were as follows. For the "is it animate?" task, the 20 similar prime-target pairs were birds and mammals; the 20 dissimilar primes were tools and clothing, the similar primes from the made by humans task. For "does it grow?", the 20 similar prime-target pairs were fruits and vegetables; the 20 dissimilar primes were furniture, kitchen items, vehicles, and weapons, primes from the object task. For "is it made by humans?", the 20 similar prime-target pairs were tools and clothing; the 20 dissimilar primes were birds and mammals, primes from the animate task. For "is it an object?", the 40 similar prime-target pairs were furniture, kitchen items, vehicles, and weapons; the 40 dissimilar primes were birds and mammals, primes from the animate task.

Table 1 shows how fillers were constructed. Each example prime-target pair in Table 1 corresponded to 10 trials in the “is it animate?”, “does it grow?”, and “is it made by humans?” tasks, 20 trials in the “is it an object?” task. Take the "is it animate?" task as an example. There were 50% animate and 50% inanimate targets. In a priming task in which two words are presented in succession, the prime can potentially cue the identity of the target, biasing a subject's response. Therefore, given an animate (or inanimate) prime, the probability of an animate target was 0.5. The relatedness proportion, or proportion of similar to dissimilar trials, was 0.25 (20 of 80 trials involved similar prime-target pairs). The examples in bold text corresponded to similar and dissimilar test trials; all others were filler trials. Additionally, there were 8 lead-in practice items. A separate practice session was administered in which subjects indicated whether or not a word was an adjective. There were 24 practice trials in total, 3 trials per Table 1 example.

Table 1 about here

Two lists were constructed for each task. Each list contained half of the targets with similar primes, and the opposing half with dissimilar primes. A subject saw no test or filler stimuli more than once, except the animate primes (birds and mammals), which a subject saw twice. Half of the animate primes appeared as similar primes in the animate task; the opposing half appeared as

dissimilar primes in the made by humans task, and all of them appeared as dissimilar primes in the object task.

A recognition test was administered following the practice session and each task to encourage subjects to read the primes. Five primes were used as recognition test stimuli and 5 words not appearing in any task were used as foils.

Procedure. Each subject performed all 4 tasks; they were blocked. Because 24 (4!) task orderings were possible, 24 subjects were randomly assigned to list 1 and 24 to list 2. A subject was told that all words were meant to be nouns. Each priming trial proceeded as follows: an intertrial interval of 1500 ms; an asterisk in the center of the screen for 250 ms; a 250 ms pause; a prime for 200 ms; a mask, &&&&&&&&&&, for 50 ms; a target until a response was made. The experimenter informed each subject about the sequence of a trial, asking him to silently read the first word (the prime), then respond as quickly and accurately as possible to the second word (the target). Examples were given to portray what was meant by: animate (something that is alive, a cockroach is but a keyboard is not); object (something that is tangible or concrete, a pair of scissors is but the sky is not); made by humans (something that is manufactured by people, a razor is but a butterfly is not); and grows (something that grows on its own, a spider grows but a door does not). Subjects performed an "is it an adjective" practice session. For the recognition tests that followed the practice session and each task, a subject was asked to indicate whether or not each of 10 items appeared as a prime in the previous task. It took approximately 40 minutes for a subject to complete the experiment.

Design. There were 2 factors: task, with 4 levels - animate, object, made by humans, grow; and prime type, with 2 levels - similar, dissimilar. Analyses of variance were performed using subjects and items as random variables. In the analysis by subjects, both task and prime were within subjects factors. In the item analysis, because a target was used as a unit of analysis, prime type was a within items factor (the same target appeared with both types of primes), and task was a between items factor (the targets for the 4 tasks were different). The regression analyses are described below.

Results

Analyses of variance. Latency and accuracy of subjects' responses were recorded. Errors occurred on less than 4% of the trials for each task: animate, 3.9%; object, 2.4%; grow, 3.1%; and made by humans, 3.0%. Subjects made errors on 2.2% of similar prime-target trials and 3.7% of dissimilar prime-target trials. Due to their low frequency, they were not further analyzed. All response latencies greater than 3 standard deviations from the mean latency for experimental targets in each task were replaced by the cutoff value. The percentage of response latencies affected by this procedure were: animate, 1.4%; object, 1.5%; grow, 0.6%; and made by humans, 2.4%.

Mean response latency (with standard error) by task and prime type are displayed in Table 2. Analyses of variance were conducted using response latency as the dependent variable. No interaction between task and prime type was apparent, $F < 1$ by subjects, $F(3, 96) = 1.09$, $MS_e = 3184.71$, $p > .4$ by items. A significant priming effect was found; subjects responded to a target more quickly when it was preceded by a similar prime (767 ms) than by a dissimilar prime (810 ms), $F(1, 47) = 36.02$, $MS_e = 4961.59$, $p < .0001$ by subjects, and $F(1, 96) = 31.98$, $MS_e = 3184.71$, $p < .0001$ by items. A significant effect of task was found, $F(3, 141) = 47.77$, $MS_e = 17759.21$, $p < .0001$ by subjects, and $F(3, 96) = 38.55$, $MS_e = 9993.29$, $p < .0001$ by items. Mean response latency (and standard error) by task was: animate, 723 ms (14 ms); grow, 758 ms (9 ms); object, 752 ms (10 ms); and made by humans, 926 ms (17 ms).

Table 2 about here

Planned comparisons were performed to investigate priming effects in each task (see Table 2 for means). Subjects were faster to make decisions to targets preceded by similar primes in 3 of the 4 tasks: "is it animate?", $F(1, 188) = 13.10$, $MS_e = 5541.14$, $p < .01$ by subjects, and $F(1, 96) = 12.46$, $MS_e = 3184.71$, $p < .01$ by items; "is it an object?", $F(1, 188) = 9.98$, $MS_e = 5541.14$, $p < .01$ by subjects, and $F(1, 96) = 15.70$, $MS_e = 3184.71$, $p < .01$ by items; and "is it made by humans?", $F(1, 188) = 6.93$, $MS_e = 5541.14$, $p < .01$ by subjects, and $F(1, 96) = 8.82$, $MS_e = 3184.71$, $p < .01$ by items. However, the priming effect in the "does it grow?" task was not significant, $F(1, 188) = 3.40$, $MS_e = 5541.14$, $.5 < p < .1$ by subjects, and $F(1, 96) = 1.66$, $MS_e = 3184.71$ $p > .1$ by items.

Regression analyses. Two items, EMU-OSTRICH and STARLING-CROW, were removed because it was apparent that few people knew the meanings of EMU or STARLING. In a similarity rating task described in McRae (1991), 11 of 20 subjects reported being completely unfamiliar with EMU, and 8 were unfamiliar with STARLING. Further evidence came from Experiment 1, in which 13 of 30 subjects listed either incorrect properties for EMU (e.g., *has fur*) or no properties at all (10 for STARLING). Finally, in familiarity norms in which 1 represented completely unfamiliar and 7 represented extremely familiar, EMU received a mean familiarity rating of 2.45 (190th of 190) and STARLING received 2.95 (184th of 190).

Similarity among the 88 pairs of concepts was computed using both the individual and correlated properties representations. Thus, prime-target concept similarity was used to predict mean priming effect on an item-wise basis. Similarity for the dissimilar prime-target pairs was not included in regression analyses because it carried no information; 74 of 88 dissimilar prime-target pairs shared no properties, 13 shared 1, and 1 shared 2.

Concept similarity, as measured by overlap of individual properties, was used to predict priming effects on an item basis by predicting mean response latencies for similar prime-target pairs (LAMP-CHANDELIER) after mean response latency for dissimilar pairs (GOOSE-CHANDELIER) had been entered into a regression equation. Because all probabilistic theories of conceptual structure assume that individual properties are encoded in conceptual memory, similarity in terms of correlated properties was used to predict residual priming effects after similarity in terms of individual properties had been entered into the equation. Analyses were conducted separately for artifacts and biological kinds. Overlap of individual properties was a significant predictor of priming effects for artifact prime-target pairs, $R^2 = .15$, $F(1, 51) = 8.963$, $p < .005$, but not for biological kinds, $R^2 = .04$, $F(1, 31) = 1.297$, $p > .3$. Conversely, overlap of correlated property pairs significantly predicted priming effects for biological kinds, $R^2 = .25$, $F(1, 30) = 10.190$, $p < .004$, but not for artifacts, $R^2 = .003$, $F < 1$.

In order to better understand why the correlated properties representation predicted biological kind priming effects, we manipulated one aspect of the representation at a time. There were four major characteristics of the correlated properties representation. First, it contained only correlated property pairs, rather than all pairs. Second, property pair violations were represented. Third, only properties that occurred in at least three concepts were included. Fourth, property pairs were not weighted in the representation according to strength of correlation.

Importantly, prediction of priming effects failed when the representation included all property pairs, whether or not they were correlated, $R^2 = .01$, $F < 1$. Predictions were also nonsignificant when information about violated correlations was removed, $R^2 = .09$, $F(1, 31) = 2.932$, $p > .09$. When properties that occurred in two concepts were included in the analyses, predictions were not affected, $R^2 = .23$, $F(1, 31) = 9.040$, $p < .006$. However, if unique properties were included, predictions were nonsignificant, $R^2 = .04$, $F(1, 31) = 1.379$, $p > .2$, probably because if a property occurred only once in the norms, many correlations involving it were probably spurious. Finally, weighting pairs by how strongly they were correlated might increase predictive accuracy. To construct such a representation, significantly correlated pairs were represented by one to ten dimensions, contingent upon amount of shared variance. Strength of correlation was coded

discretely by using it to determine the number of units assigned to a property pair using the following rule (number of units shown in parentheses): (1) .064 < R^2 < .164; (2) .164 < R^2 < .264; (3) .264 < R^2 < .364; and so on. Predictions using this scheme were similar to the unweighted version of the correlated properties representation, R^2 = .21, $F(1, 31) = 8.004$, $p < .009$. Finally, note that all of these representations predicted less than 1% of the variance of artifact priming effects after similarity in terms of individual properties was entered into the equation.

In summary, concept similarity in terms of individual properties predicted priming effects for artifacts whereas concept similarity in terms of correlated property pairs predicted priming effects for biological kinds. The important elements of the correlated properties representation were: including only correlated pairs; including information about violations; and excluding unique properties. Weighting property pairs by correlational magnitude did not affect predictions.

Discussion

Experiment 3 demonstrated that property intercorrelations are encoded in conceptual memory and influence the computation of concepts from words. Concept similarity in terms of correlated properties predicted biological kind short interval priming effects, although it was not required for artifacts. That is, for biological kinds but not artifacts, similarity in terms of correlated properties predicted priming effects and similarity in terms of individual properties failed. Thus, the nonlinearities introduced by correlated properties were required to properly measure biological kind similarity. There are two sources of nonlinearity introduced by correlated properties (as discussed in Section 3.0). First, similarity in terms of correlated property pairs is a positively accelerated function when plotted against shared individual properties. The more densely intercorrelated are the properties of concepts, the more nonlinear the function. It was shown in Section 3.0 that biological kinds have more densely intercorrelated properties than do artifacts. Second, concept similarity in terms of correlated properties was nonlinear in that it could potentially decrease with number of shared properties. If concept1 possessed both properties of a correlated pair, but concept2 possessed only one (i.e., violated the correlation), then they were less similar than if concept2 possessed neither property. Thus, similarity over individual properties was incorrect to the extent that these cases existed. The number of cases where one concept possessed both properties of a pair and the other possessed only one was computed for each of the 88 prime-target pairs used in the Experiment 3 regression analyses. The mean number of these cases was greater for biological kind pairs (mean = 21.3, se = 1.7) than for artifacts (mean = 8.4, se = 1.1), $t(86) = 6.819$, $p < .0001$. In summary, due to the density of property intercorrelations and violations in biological kinds, concept similarity is nonlinear to the extent that it cannot be adequately captured in terms of individual properties (i.e., independent cues). In contrast, because the properties of artifacts tend to be sparsely intercorrelated, individual properties are sufficient to capture their similarity.

In Section 3.1, it was claimed that short interval semantic priming is a sensitive measure of concept similarity. The assumption seems valid given the strong effect of individual and correlated properties that was obtained. However, if priming is sensitive to concept similarity in terms of property overlap, perhaps more than 25% of the variance should have been explained. Two factors may have impaired predictions. First, although it was shown that property-based concept similarity is a basis of short interval semantic priming effects, previous studies have shown priming due to other types of semantic relationships such as co-occurrence of entities (e.g., BREAD and BUTTER,) or of words in speech or text (e.g., PRIVATE and PROPERTY, Hodgson, 1991). Predictive accuracy of property-based similarity would have been reduced to the extent that these relationships existed between the primes and targets of Experiment 3.

Second, taking the "is it animate?" task as an example, it can be seen from Table 1 that the primes and targets for the similar test trials were animate (e.g., CARIBOU-MOOSE), but the primes for the dissimilar control trials were inanimate (e.g., DRILL-MOOSE). Subjects may have making implicit semantic decisions to the primes. If so, the decision stage may have been speeded for a target when preceded by a prime that was animate (or an object, or grew, or man made) as

compared to when it was preceded by a prime that was inanimate (or not an object, or not grown, or not made by humans). This cuing of the decision would not have been sensitive to subtle manipulations of similarity between prime and target; only the animacy of the prime would have mattered. Therefore, concept similarity would not have accounted for this variance in the regression analyses, resulting in a decrease in the proportion of variance accounted for without altering the pattern of results.

Relationship to previous results. McRae (1991) found that concept similarity in terms of correlated properties was a better predictor of performance for biological kinds than for artifacts in a same/different category decision task, although results were not as clear. Furthermore, in both Experiment 3 and McRae's same/different task, individual properties better predicted performance for artifacts but not for biological kinds. In contrast, some researchers have previously failed to find a significant contribution of correlated properties when predicting typicality (Malt and Smith, 1984; Murphy & Wisniewski, 1989) and similarity ratings (McRae, 1991). Two critical differences exist between rating typicality or similarity and the unconscious aspects of concept similarity tapped by the short interval priming task. First, a priming task allows for sensitive analyses because it promotes a design in which a target can serve as its own control; however, this is not true of typicality or similarity judgments. Second, typicality and similarity ratings are slow, conscious, deliberate tasks, and the information on which subjects base their judgments for these tasks is unclear. For example, it may be that judgments are based on a limited number of dimensions (Rips, Shoben, & Smith, 1973), or, perhaps they are dominated by the types of knowledge that knowledge-based theories emphasize (McRae, 1991; Rips, 1989). In contrast to conscious judgment tasks, short interval semantic priming effects can be seen as reflecting concept similarity at a more basic representational level because effects depend on rapid successive computation of two concepts. A mechanistic account of similarity-based priming is provided in Section 6, where a simulation of this experiment is presented.

Semantic relatedness. In Experiment 3, the magnitude of short interval priming effects was predicted by concept similarity, as measured by overlap of individual and correlated properties. This experiment represents the first demonstration in which the size of priming effects were predicted on an item basis. There is a vast empirical literature on semantic priming effects and several models have been proposed (e.g., Becker, 1980; Collins & Loftus, 1975; McNamara, 1992; Neely & Keefe, 1989; Posner & Snyder, 1975; Ratcliff & McKoon, 1988). Short interval semantic priming has been explained by reference to an automatic spreading activation process internal to lexical memory (memory for word-specific information). According to spreading activation theory, when a prime is recognized, its lexical node is activated and activation spreads to semantically related items nearby in the lexicon. Facilitation results because a target's lexical node is partially activated prior to its presentation. There are, however, numerous conceptions of the informational basis of semantic relatedness. This confusion is underscored by a recent article by Hodgson (1991), in which he compared six types of semantically related prime-target pairs: synonyms (e.g., TOLERATE-ALLOW), antonyms (e.g., ENEMY-FRIEND), within-category exemplars (e.g., MIST-RAIN), category-exemplar pairs (e.g., HOUSE-COTTAGE), conceptual associates (e.g., DOVE-PEACE), and phrasal associates (e.g., PRIVATE-PROPERTY). Although priming effects were obtained using the lexical decision task, Hodgson found no differences among these stimulus types in terms of the magnitude of priming effects. He concluded that priming is caused by a retrospective mechanism, as described by Neely and Keefe (1989). Other attempts to investigate the informational basis of semantic relatedness have also been unsuccessful (e.g., Lewis, 1970; Warren, 1977).

Effects of concept similarity have been obtained in other priming studies, although they have not been portrayed as such. den Heyer et al. (1983), Favreau and Segalowitz (1983), Neely (1977), and Smith, Briand, Klein, and den Heyer (1987) found priming in the lexical decision task using category name-exemplar name pairs. Whether a superordinate concept is stored as a prototype or computed from exemplars on-line, category concepts are assumed to be similar to typical exemplar concepts (e.g., Hintzman, 1986; Rosch & Mervis, 1975). Given an account of

short interval semantic priming founded on property-based concept similarity, category-exemplar priming effects would be expected. Furthermore, four of the six types of prime-target pairs that Hodgson (1991) used may reflect concept similarity: synonyms, antonyms, pairs of exemplars chosen from within a category, and category-exemplar pairs.

The remaining types of prime-target pairs used by Hodgson (1991), conceptual associates (e.g., DOVE-PEACE) and phrasal associates (e.g., PRIVATE-PROPERTY), are semantically related through association rather than similarity. Semantic relatedness has been defined in this way in a number of theories of short interval semantic priming (see Neely, 1990 for a review; McNamara, 1992). The source of associative strength is seldom defined. It may be due to temporal contiguity in speech or text (McKoon & Ratcliff, 1992) or co-occurrence within a proposition (McNamara). Regardless, associative strength has been operationalized using free association norms. It is assumed that the same mechanism is responsible for both free association and short interval semantic priming effects. To evaluate whether the main effect of conceptual similarity in Experiment 3 could be explained by associative relatedness, a subset of the stimuli were selected that were semantically similar, but not associatively related. Twenty subjects were read the 100 primes and were asked to respond with the first thing that came to mind. A set of 65 semantically similar prime-target pairs that were not associatively related were chosen by retaining only those pairs for which 0 or 1 out of 20 subjects responded with the target when free associating to the prime. Analyses of priming effects were conducted using these items and collapsing across task. With associative relatedness removed, a 39 ms priming effect remained: similar pairs (mean = 784 ms, standard error = 14 ms); dissimilar pairs (mean = 823 ms, standard error = 15 ms); $t(64) = 3.98$, $p < .0001$ by items. Using a similar method, Fischler (1977) has also found a short interval priming effect with semantically similar, but not associatively related, prime-target pairs.

McNamara (1992) has claimed that an analysis such as this is insufficient because it fails to take into account mediated associations (e.g., bracelet to wrist to watch, McNamara & Altarriba, 1988). There are two reasons why the priming effects found in Experiment 3 cannot be attributed to mediated priming. First, the effects are much larger than those found in mediated priming studies (39 ms versus 10-20 ms). Second, we did not simply demonstrate an overall priming effect using prime-target pairs corresponding to similar concepts. Rather, the magnitude of priming effects was predicted by property-based concept similarity. This method provides a strong argument that the priming effects found in Experiment 3 resulted from concept similarity.

3.3 Summary of Sections 2 and 3

Experiments 2 and 3 demonstrated that humans encode correlations among properties of entities in conceptual memory. Analyses of the production norms collected in Experiment 1 showed that the properties of biological kinds are more densely correlated than those of artifacts. In line with this result, Experiment 3 showed that property intercorrelations were more important for predicting priming effects for biological kinds than for artifacts. It also demonstrated that short interval semantic priming effects can be predicted by concept similarity in terms of individual and correlated properties.

4.0 A MODEL: COMPUTING WORD MEANING FROM WORD FORM

The stimuli of Experiments 1, 2, and 3, as well as of most studies of real-world concepts, were words. Therefore, results suggest that correlated properties play a role in the computation of word meaning. In this section, we outline a connectionist model that is based on this principle. The model is then used to simulate the results of the property verification task (section 5) and the short interval semantic priming task (section 6). Furthermore, in section 7, we empirically test the model's prediction that computing conceptual representations from words depends on stored knowledge of property intercorrelations.

Connectionist models are a natural way to capture the notion that stored property intercorrelations are central to conceptual memory. For example, an iterative connectionist model

with a correlational learning rule (e.g., a Hopfield network, 1982, 1984) can naturally be used to encode correlations among the properties of entities in its environment. Furthermore, these models may be used to shed light on the processes by which property intercorrelations influence people's performance on semantic tasks. The goal of our modeling was to use an explicit computational model to investigate a theory in which encoded knowledge of the correlational structure of semantic space plays a key role in the computation of word meaning. To achieve this goal, we implemented a modified Hopfield network that used encoded correlations among properties to construct basins of attraction that were fundamental to its computational dynamics.

There is independent motivation for hypothesizing that encoded property intercorrelations play an important role in the computation of word meaning. The mapping from word form to meaning contains few regularities, particularly at the monomorphemic level; thus, words that sound or are spelled similarly do **not** tend to have similar meanings (e.g., MAT, RAT, THAT, HAT, CAT, and SAT). As Hinton and Shallice (1991) have pointed out, it is difficult for a network with distributed representations to learn a system in which similar inputs (e.g., the orthographic or phonological forms of the words listed above) must be mapped onto dissimilar outputs (e.g., their meanings). This observation led them to suggest that the computation of word meaning from orthography (or phonology) may require two overlapping stages. First, a computation from word form puts the system into an appropriate region of conceptual space. The second stage uses relationships among semantic properties to clean up the pattern, or to drive the system to a stable state that correctly corresponds to the concept in question. In a series of simulations, Hinton and Shallice (1991) and Plaut (1991) have shown that basins of attractions that are formed by encoding the higher order structure of semantic space are critical for simulating and explaining the word naming behavior of deep dyslexic patients. Hinton and Shallice (1991) used an iterative backpropagation network for their investigations of deep dyslexia. Their model included a standard feedforward path from orthography to conceptual memory through a set of hidden units. The iterative component of their network consisted of a loop from the conceptual units through a set of hidden units back to conceptual memory. Basins of attraction were formed in this part of the network.

Our interest in exploring implications of encoding property co-occurrence information in conceptual memory led to using a different type of network than used by Hinton and Shallice (1991), who were primarily interested in studying deep dyslexia. We used a Hopfield (1982, 1984) network, in which learning is accomplished using the Hebb (1949) correlational learning rule. In a Hopfield network, if a semantic unit corresponds to a property, the encoding of property intercorrelations is transparent and easily analyzed. In contrast, in Hinton and Shallice's network, the semantic loop encoded the correlational structure of the learned entities in a highly opaque manner. A second advantage of a Hopfield network is mathematical/analytical; its operation can be described by a well-defined energy function. In addition to these advantages, like iterative backpropagation networks, computations in Hopfield networks are dependent on attractors constructed through encoded relationships among the units (properties). Finally, a major advantage of our model, unrelated to the type of network used, was that our conceptual representations were based on people's actual descriptions of entities; they were based on the norms collected in Experiment 1. In contrast, in previous models, concepts were represented either by sets of properties created by the researchers themselves (e.g., Hinton & Shallice) or by random patterns of activation (e.g., Farah & McClelland, 1991; Masson, 1992; Sharkey, 1989).

4.1 Architecture

Figure 3 shows the model's architecture. The input was a distributed representation of word form. To derive it, we computed all triples of letters from the words that were included in the model, where spaces at the beginning and end of a word were treated as characters. This resulted in 379 word form units and provided a sparse distributed representation that roughly preserved item similarity. Because the human studies used written words, we sometimes refer to the input units as coding orthography. However, the critical point was that no systematic relationship

existed between the input patterns and conceptual representations. We assume that, when all words are represented, the lack of regularity in the mappings from orthography or phonology to conceptual memory results in a computational system in which concepts cannot be computed directly from orthography or phonology. Due to this characteristic, the fact that encoded knowledge of property co-occurrences are available to influence a concept's convergence is critical to computing word meaning.

Figure 3 about here

The output was a representation of word meaning. Each of the 646 output units corresponded to a semantic property from the production norms (properties were represented locally). A concept was a distributed representation over the property units. In this scheme then, DOG and TRUCK corresponded to nonoverlapping conceptual patterns - nonoverlapping because DOG and TRUCK shared no semantic properties. The representation was binary (0 = property is absent, 1 = property is present); no production frequency information was included. From among the 190 normed entities, 84 concepts corresponding to 42 prime-target pairs were included in the model. Fourteen items from the property verification task were also included.

The conceptual/output units were fully bidirectionally interconnected to allow the entire range of relationships among semantic properties to be encoded (not every connection is shown in Figure 3). The orthographic/input units were fully unidirectionally connected to the conceptual units, and were not interconnected.

4.2 Determination of Connection Weights

The model instantiated a theory in which it was assumed that people encode the correlational structure of their environment during their experiences with the world. People also learn to comprehend speech and text; part of this process involves learning to map orthographic and phonological patterns onto conceptual ones. The model's training corresponded to this theory in that the weights among the property units were determined independently of the weights connecting orthographic to conceptual units. Training for each section was accomplished using a single batch matrix multiplication, rather than iteratively training on each pattern. Theoretically, the methods are similar, with matrix multiplication being pragmatically preferable in terms of computational considerations. Note that all entities were equally familiar to the model because each concept was used once when the weights were determined.

The weights between conceptual units were set using the Hopfield (1982, 1984) learning rule slightly modified to be optimal for sparse patterns (Tsodyks & Feigelman, 1988). The learning rule for the conceptual units was:

$$\underline{w}_{ij} = 1/\underline{n}_1 \sum_p [(\underline{x}_{ip}-\underline{m}_p)(\underline{x}_{jp}-\underline{m}_p)] \qquad (1)$$

where $\underline{w}_{ij}$ represented the weight of the connection between $unit_i$ and $unit_j$ (connections were symmetric, $\underline{w}_{ij} = \underline{w}_{ji}$), $\underline{n}_1$ was the number of conceptual units, and $\underline{x}_{ip}$ represented the absence (0) or presence (1) of $property_i$ in $concept_p$. The term $\underline{m}_p$ is the proportion of properties in $concept_p$ (number of properties in $concept_p$ divided by 646); this term made the learning rule sensitive to concept sparsity. Its effect can be demonstrated using the semantic pattern corresponding to EAGLE, for which 17 of 646 possible properties were present, a proportion of .0263. If two EAGLE property units were active (e.g., *has wings* and *flies*), their connection was strengthened by 0.9481. If two units were inactive (e.g., *made of wood* and *has a handle*), their connection

was strengthened by 0.0007. Finally, a connection between an active and inactive unit (e.g., *has wings* and *has a handle*) was strengthened in the negative direction by an intermediate quantity, 0.0256.

In addition to allowing greater storage capacity, the modified Hopfield rule for sparse patterns (Tsodyks & Feigelman, 1988) captured the following intuition. Because the number of properties possessed by an entity in relation to the set of all possible properties in the world is minute, the absence of a property generally carries little or no information. For example, when someone attends to a dog, the person notices the properties that it possesses, but does not tend to notice ones that the dog does not possess, such as that it does not have a handle. The learning algorithm for sparse patterns appropriately captures the relative importance of the co-occurrence of two properties; weights are strengthened between co-occurring properties to a greater extent than between two absent ones. It should be noted that the situation described above is qualitatively different from one in which a dog is missing a typical DOG characteristic; for example, if a dog *has three legs*, people do tend to notice (and so would the model).

The connections between input and output units were computed using a similar Hebbian learning rule (Hebb, 1949). Specifically, the orthographic to conceptual learning rule was:

$$z_{ij} = 1/n_2 \sum_p [(a_{ip})(x_{jp}-m_p)] \qquad (2)$$

where a_{ip} was the normalized input pattern at input unit$_i$, x_{ip} and m_p were the same as in Equation 1, and n_2 was the number of input units. An input pattern was represented by a 1 at each unit corresponding a letter triple contained in the word, and 0's elsewhere. Patterns were normalized to remove word length bias according to the following equation:

$$\sum a_{ip}^2 = 4 \qquad \text{for all } p \qquad (3)$$

4.3 Computing a Word's Conceptual Representation

The dynamics of unit activations are given by the following equation:

$$x_i(t+1) = g(c_1 \sum_j (z_{ji} * a_i) + c_2 \sum_j (w_{ji} * x_j(t)) - \emptyset) \qquad (4)$$

$$\text{where } g(x) = 0.5 \tanh(c_3 x) + 0.5$$

The value $x_j(t)$ was the activation of unit$_i$ at time t; x, a, z, and w were defined above. Synchronous updating was used for speed and ease of simulation; we would not expect qualitative differences using asynchronous updating. The following values were used for the simulations presented below, although similar results were obtained with various combinations of ranges of values for the constants: $c_1 = 0.85$, $c_2 = 0.33$, $\emptyset = 0.0105$, $c_3 = 400$, and g was a steep sigmoid. The algorithm was not overly sensitive to the steepness of the sigmoid (given by c_3) but a step function did not work (see below).

In the computation of a word's meaning, its orthographic (word form) representation was used as input. Thus, it was assumed that preliminary analyses had resulted in an internalized distributed representation of how the word was spelled. The network was initialized by setting 60 randomly chosen conceptual units to 0.25 because concepts activated approximately an average of 15 units. When 60 units were set at 0.25, the total activation in the system was similar to its normal stable state. Given that activation began as random and close to zero, in the first iteration (or time step), activation of conceptual units resulted primarily from orthography. Clamping an orthographic pattern moved the system in semantic state space to a point within the appropriate

basin of attraction. On further iterations, updates based on orthographic and conceptual units resulted in gradient descent in state space toward the bottom of the attractor basin, a low energy point in state space corresponding to a known concept. For this reason, the sigmoidal function performs better than the step function. Because it allows concept unit activations to increase slowly, it helps to increase the activation of other units. Thus, convergence latency depended on both knowledge of the mapping from word form to meaning and encoded property intercorrelation information stored in the weights between conceptual units.

The model can be understood as a Hopfield (1982, 1984) network with the effect of the orthographic connections as thresholds that were variable over units and patterns. Thus, the semantic interconnections determined the general topology of the energy function---which combinations of properties were stable based on its experience with property co-occurrence. The orthographic connections decreased the energy of states involving semantic properties that occurred in words with activated word form units. They also provided an initial state to the network.

In summary, the model embodied two important principles: it naturally encoded how properties co-occurred in entities; and it used this encoded knowledge of conceptual structure to drive the system to a stable state.

4.4 Encoding Correlations

The weights encoded information that was comparable to the Pearson product moment correlation computed for the correlated properties representation. As described above, due to sparsity of the concepts, in the determination of the weights, the largest contributing factor was property co-occurrence. This factor also dominated correlations between property pairs. Pearson correlation directly varies with the sum of the products of the standard scores on the two variables. Because property vectors were sparse, the mean of the production frequencies was less than 0.4 for all properties, and less than 0.1 for 96% of them. Therefore, a production frequency of 0 fell close to the mean, so that an absent property contributed little to the computation of the correlation coefficient relative to one that was present (production frequency between 5 and 30). The correlated properties representation differed from the model in that, when computing correlations between property pairs, only properties listed for three of more concepts were used. In contrast, all properties were interconnected in the model. In summary, the model and the correlated properties representation corresponded in that they both were sensitive to the sparsity of semantic space.

4.5 Performance Analyses

The model was trained on 84 concepts spanning the 10 categories. There were 10 BIRDS, 10 MAMMALS, 8 FRUIT, 10 VEGETABLES, 8 CLOTHING, 6 FURNITURE, 8 KITCHEN ITEMS, 8 TOOLS, 8 VEHICLES, and 8 WEAPONS. Concepts were chosen to facilitate simulations of Experiments 2 and 3; 14 concept pairs from the property verification task were included, as were 42 similar prime-target pairs.

The network successfully learned 82 of the 84 concepts; this was approximately the maximum number of concepts that could be learned by the network due to crosstalk among concept vectors (for discussions of storage limitations in Hopfield networks, see Hertz, Krogh, & Palmer, 1991; Hopfield, 1982, 1984). After training, we tested the network's retention of each word by randomly initializing 60 conceptual units to a value of 0.25. The orthographic representation of a word was then clamped and the model iterated ten times. Two performance measures were recorded after each iteration: (1) error, the sum of squared error between the computed concept vector and the target vector; and, (2) activated properties, where a unit was considered to be on if its activation was greater than 0.5 (target values for units were either 0 or 1). To avoid peculiarities due to a particular set of initial values, results were averaged over five runs, each of which used a new set of randomly chosen initial values for the conceptual units.

Figure 4 shows a graph of the mean sum of squared error for the 84 concepts, as well as for CRAYON, the most slowly converging concept, and JET, a concept that converged, then activated

a number of incorrect properties. On average, error steadily decreased until time step five, at which time it stabilized. That is, concepts tended to stabilize after five time steps, at which point the mean error was 0.623. This error was quite low; by comparison, if the network simply turned all units off, which might be a reasonable solution given the sparsity of the patterns, the average expected error would have been 16.2. CRAYON and JET were interesting and contrasting cases. Because the properties of CRAYON were sparsely intercorrelated, it was slow to converge. JET, on the other hand, contained properties that were both densely intercorrelated and correlated with a number of bird-like properties. Therefore, JET converged quickly, but, error then increased because a number of bird-like properties that were strongly intercorrelated with other properties of JET became activated (*is a bird*, *has feathers*, *has a beak*, and *eats*).

Figure 4 about here

In general, errors made by the network followed a pattern similar to that of JET. Hopfield (1982, 1984) networks are pattern completion devices. Given that this network's behavior was primarily driven by encoded property co-occurrence, it tended to activate properties that were strongly intercorrelated with properties of the concept, but were not part of the concept according to the norms. After four iterations, a total of 10 properties from 8 concepts were erroneously activated. By seven iterations, it had increased to 29 properties from 16 concepts, and was stable through ten iterations. The 16 concepts with erroneously activated units contained densely intercorrelated properties (mean intercorrelational strength = 979, se = 127) as compared to the mean for the 84 concepts (579). Interestingly, for 14 of those 16 concepts, the model activated properties that were true of the entities, even though less than 5 of 30 subjects had listed the property in the norms. For example, according to the model, a BUDGIE *has wings*, a BUZZARD and an EAGLE *are animals*, a CHICKEN and a DUCK *are large*, a CANNON *is dangerous*, a CARROT, a RADISH, and a ZUCCHINI *are edible*, a CARROT *has leaves*, a HAWK *eats*, a MOUSE *has four legs*, a MISSILE *is loud*, and TROUSERS *are worn by women*. However, although the model believes that a CANARY is *an animal,* it also thinks that it *is loud* and *is large*. Furthermore, for both JET and CAT, bird-like properties were erroneously activated. Again, these errors illustrate that the model is completing patterns based on property intercorrelations. Because the magnitude of the intraconceptual weights are determined by the tendency for properties to occur together in the world, the pattern completion behavior of the model tends to activate properties that are actually true of an entity.

In contrast, a group of 12 concepts whose properties were sparsely intercorrelated (mean = 135, se = 31) were slow to converge. After four iterations, the activation of 62 properties distributed across 12 concepts remained erroneously below 0.5. After seven iterations, 10 properties of CRAYON remained activated below 0.5, as well as one property of YAM; after ten iterations, only the unit corresponding to YAM-*is like a potato* remained incorrectly off. Because properties of these concepts were sparsely intercorrelated, mutual activation among a concept's properties was relatively weak and activation accrued slowly.

In summary, the model learned 82 of the 84 concepts to the point where, after seven to 10 iterations, only a few property units were activated that were not part of the representation as specified by the norms. Furthermore, these additional property units were activated because they were highly correlated with other properties within the concept, and were, for the most part, appropriate.

In Sections 5 and 6, we use the model to help understand the results of Experiments 2 and 3. Mechanistic descriptions of property verification and short interval priming are presented and these explanations are explicitly simulated.

5.0 PROPERTY VERIFICATION

In this section, the model is used to guide explanations of the influence of property intercorrelations on verification. On each trial in Experiment 2, a concept (e.g., DEER) was presented for 400 ms, followed by a target property (e.g., *hunted by people*). Subjects indicated as quickly as possible whether the target property was reasonably true of the concept. With production frequency equated, target properties that were strongly intercorrelated with other properties within a concept were verified more quickly than matched targets that were weakly intercorrelated. In sections 5.1 and 5.2, two plausible explanations of the influence of correlated properties are outlined and simulated.

The model contained 14 items from Experiment 2 that differed in terms of the strength with which the target property was correlated with other properties of the concept, strong (mean = 163, se = 16), weak (mean = 37, se = 9), $t(13) = 7.595$, $p < .0001$. Note that, in sections 5, 6, and 7, we report measures derived from statistical analyses of the representations based on the norms using only the 84 concepts on which the model was trained (versus sections 2 and 3, in which the individual and correlated properties representations were based on all 190 normed entities). The groups were also roughly equated on three variables (means and standard errors in parentheses): production frequency, strong (12, 2), weak (12, 2), $t(13) = 0$; concept familiarity: strong (5.4, 0.2), weak (5.0, 0.3), $t(13) = 1.348$, $p > .2$; and number of individual properties listed per concept: strong (18, 1), weak (17, 1), $t(13) = 1.677$ $p > .1$. For this subset of 14 items, human verification latency was significantly faster for the strongly intercorrelated group (788 ms, 34 ms) than for the weakly intercorrelated group (998 ms, 51 ms), $t(13) = 7.595$, $p < .001$ by items. In the simulations, because all concepts were trained equally, familiarity was equated.

5.1 Differential Sensitivity of Norming and Verification

The first explanation of the influence of correlated properties on property verification hinges on the assumption that production frequency is less sensitive than verification latency as a measure of a property's activation by a concept name. Production frequency measures the probability that a property is one of the first few generated by a subject when reading a concept name. For example, if subjects list an average of eight properties for a concept, a property's production frequency roughly indicates the frequency (or probability) that a property is generated as one of the top eight. Recall that the norming task of Experiment 1 imposed no time limits on subjects; therefore, this measure of production frequency was insensitive to the speed with which properties were generated. In contrast, in the property verification task, a subject was asked to indicate, as quickly as possible, whether or not a specific property was reasonably true of an entity. Thus, verification latency provided a more direct, precise, and sensitive measure of the relationship between a concept and a property. It might be assumed that subtle differences existed in a target property's activation that were too small to influence production frequency, but large enough to be detected in the property verification task. These differences resulted from encoded knowledge of property co-occurrences.

In summary, it might be assumed that equating for a target property's production frequency roughly equated for its activation given a concept name. However, because production frequency was an imprecise measure of availability, a difference was found in the ease with which a target property was verified. Furthermore, this difference was predictable from the strength with which a property was correlated with other properties within a concept. Thus, when the target was correlated with a number of properties within a concept, it may have received a boost in activation.

To test this hypothesis, the orthographic form of each of the fourteen concepts was clamped and they were allowed to converge for 10 iterations. Activation of the target property was recorded after each iteration. Five independent runs were used as before. Figure 5 shows the mean activation of the fourteen target properties as the concepts converged. A one-way repeated measures analysis of variance showed that the activation of the target properties was significantly higher when a concept from the strongly intercorrelated group was being computed, $F(1, 13) = 10.277$, $p < .008$. An interaction between group and time step also obtained, $F(9, 117) = 5.057$, p

< .0001. Simple main effects revealed that target properties from concepts from the strongly intercorrelated group were not significantly more activated after time step 1, but were after time steps 2 through 5 ($p < .05$). The advantage was marginally significant throughout time steps 6 to 10 ($p < .07$). Thus, when property verification is viewed in this way, the model predicts an influence of correlated properties.

Figure 5 about here

5.2 Computing Meaning from a Property Name

There is a second explanation of the property verification effects that requires a slightly different conception of production frequency. Perhaps equating for normed production frequency was equivalent to equating for the strength with which a property was activated when a concept name was read. Under this assumption, the computation of a concept was **not** a source of any difference between the strong and weakly intercorrelated groups. It might be further assumed that, if property intercorrelations are encoded in conceptual memory, then when the representation of a property is computed, properties with which it is correlated may be activated. Thus, intercorrelations may influence verification when a target property is computed by activating more properties of a concept from the strongly intercorrelated group than of the corresponding concept from the weakly intercorrelated group.

Returning to Figure 2, *hunted by people* was paired in this study with DEER and DUCK. According to the norms, *hunted by people* is correlated with eleven other properties of DEER, but only four properties of DUCK. If property intercorrelations are encoded in conceptual memory, then, when the representation of *hunted by people* is computed, it will more strongly activate properties associated with wild mammals that have antlers and live in the woods than properties associated with water birds. This can be conceived in terms of a pattern completion process; when the representation of *hunted by people* is computed, it completes more of the pattern of DEER than of DUCK. Therefore, the verification decision is easier for DEER-*hunted by people* than for DUCK-*hunted by people*.

This explanation was implemented in the following manner. Because the model was not trained on orthographic representations corresponding to specific properties, the input units were turned off during this simulation. The first step was to clamp a concept. To simulate computation of the target property, the concept was left on and the node corresponding to the target property was activated. Three independents runs were used, in which the target was given an activation of 5, 10, or 15. Given that each property was represented as a single node, increased activation was necessary for one of 646 properties to significantly affect the dynamics of the system. The network was allowed to iterate five times (arbitrary). The target property unit was reactivated for each iteration because its name had remained on screen during the human experiment until the subject responded. Two measures of similarity between the computed pattern and the concept were recorded: error, and the number of activated properties that were possessed by the concept. It was predicted that, when activating target properties that were strongly intercorrelated with other properties of a concept, the system would remain in a state that more closely resembled the concept. We assumed that the less that the computed pattern drifted away from the concept when the target property was activated, the easier the verification decision.

Figure 6a shows mean error between a concept and the pattern that resulted from clamping the concept, then activating the target property. On average, the computed pattern drifted further away from a concept that possessed properties that were weakly intercorrelated with the target property. In a one-way repeated measures analysis of variance, there was a difference between the

mean error of the strong and weakly intercorrelated groups, $F(1, 13) = 5.637$, $p < .04$, but no interaction between group and time step, $F < 1$. Planned comparison revealed a significant difference between groups at each of the five time steps, $p < .0001$ for all comparisons.

Figure 6b shows a similar pattern using number of correctly activated units as the measured variable. Computed patterns were more similar to a concept when it possessed properties that were strongly intercorrelated with the target property, $F(1, 13) = 34.386$, $p < .0001$. Group and time step interacted, $F(4, 52) = 18.833$, $p < .0001$, but simple main effects showed a significant difference between groups at each of the five time steps, $p < .003$ for all comparisons. Thus, when property verification is viewed in this way, the model predicts an influence of correlated properties.

Figures 6a and 6b about here

In summary, two plausible explanations of the property verification results were presented and simulated. When a concept is computed from its name, properties that are strongly intercorrelated with others of that concept receive a boost in activation. Furthermore, when a property is computed from its name, it activates correlated properties and the resulting pattern is more similar to a concept that possesses many properties correlated with it. The model suggests that both of these factors may have been operating in the property verification task of Experiment 2.

6.0 SIMULATING SHORT INTERVAL SEMANTIC PRIMING

In Experiment 3, concept similarity in terms of individual properties predicted priming effects for artifacts, but similarity in terms of correlated properties predicted priming effects for biological kinds. Short interval priming can be understood in terms of inter-concept distance in semantic space (for similar proposals, see Sharkey, 1989; and Masson, 1992). Computing the meaning of a word can be regarded as driving the semantic system from its present state to the desired end state corresponding to the word's meaning. In a short interval semantic priming task, the prime determines the start state for the computation of the target. Because computing the meaning of a target word involves driving the semantic system from the meaning of the prime to the meaning of the target, prime-target similarity determines amount of facilitation.

Simulating short interval semantic priming is straightforward in a system that is sensitive to semantic similarity and possesses temporal qualities. Because a word's meaning is assumed to be computed in 100 to 300 ms (Gough & Cosky, 1977; Rayner, 1978), the 250 ms stimulus onset asynchrony in Experiment 3 should have been sufficient to allow a stable conceptual pattern to be computed for the prime. Therefore, for the priming simulations, the prime's orthographic pattern was clamped and its conceptual representation was computed. With the conceptual representation of the prime active, the orthographic representation of the target was then clamped. Convergence latency was recorded for the target concept.

The basic idea of this simulation was to determine if the factors that influenced priming in humans also influenced priming in the model. Specifically, it was important to show the interaction between type of representation (individual versus correlated properties) and type of entity (artifact versus biological kind). To accomplish this, we conducted the same regression analyses as in Experiment 3 with two key changes: the dependent variable was the number of iterations for a concept to converge, rather than human decision latency; and the individual and correlated properties representations were based on the 84 entities on which the model was trained, rather than all 190 normed entities. Predictions were again derived from representations based on individual properties and correlated property pairs. Linear regression analyses were used to predict

priming effects for artifacts and biological kinds from concept similarity in terms of individual and correlated properties. Because determining when a concept converges is somewhat arbitrary, three measures were used: when error dropped below 1.0, when error was within 0.1 of its value when the concept stabilized, and when error was with 0.01 of the stabilization point. It was assumed that convergence latency for a concept in the model was monotonically related to the time required for people to compute a word's meaning, and was therefore monotonically related to how quickly and easily a person could answer a question concerning its meaning. Analogous to Experiment 3, convergence latency for a target when preceded by a dissimilar prime was the first factor entered into the regression equation. In the simulation, this measure was calculated by averaging over three runs in which a concept that shared no properties with the target was used as a prime. Only 40 of the 42 priming pairs were used in the analyses because, as described above, the model did not learn to compute the meanings of CAT and JET.

The model produced the same basic pattern of results as was found in Experiment 3. Concept similarity as measured by overlap of individual properties significantly predicted priming effects for artifacts for each of the three convergence measures: less than 1: $R^2 = .27$, $F(1, 19) = 7.065$, $p < .02$; within 0.1: $R^2 = .59$, $F(1, 19) = 27.074$, $p < .001$; within 0.01: $R^2 = .58$, $F(1, 19) = 26.166$, $p < .001$. Individual properties did not predict biological kind priming effects: less than 1: $R^2 = .07$, $F(1, 15) = 1.138$, $p > .3$; within 0.1: $R^2 = .04$, $F < 1$; within 0.01: $R^2 = .01$, $F < 1$. In contrast, after similarity in terms of individual properties had been entered into the equation, similarity in terms of correlated property pairs predicted priming effects for biological kinds: less than 1: $R^2 = .09$, $F(1, 14) = 1.389$, $p > .2$; within 0.1: $R^2 = .26$, $F(1, 14) = 4.864$, $p < .05$; within 0.01: $R^2 = .25$, $F(1, 14) = 4.689$, $p < .05$; but not for artifacts: less than 1: $R^2 = .02$, $F < 1$; within 0.1: $R^2 = .14$ in the wrong direction; within 0.01: $R^2 = .00$, $F < 1$.

As in Experiment 3, we explored the characteristics of the correlated properties representation. With humans, the important aspects of the correlated properties representation were: including only correlated pairs rather than all possible pairs; including information about correlational violations; and excluding unique properties. Furthermore, weighting a property pair by its correlational strength did not influence predictions. First, as with humans, including only correlated pairs was important for predicting biological kind priming effects in the model. When all property pairs were included, predictions were nonsignificant: less than 1: $R^2 = .01$, $F< 1$; within 0.1: $R^2 = .02$, $F < 1$; within 0.01: $R^2 = .02$, $F < 1$.

The individual and correlated properties representations used for predicting priming in the model contained all of the model's knowledge about the entities. In contrast, human priming effects were predicted using only a small subset of the knowledge that people possess about the world. Therefore, some differences in the factors that are important to predicting performance might be expected. For example, in the model's case, veridical data existed about the strength of property pair correlation. Thus, it might be expected that weighting property pairs in a correlated properties representation would improve predictions of model priming effects. Consistent with this hypothesis, predictions of biological kind priming effects were more accurate when correlated pairs were weighted as described in Experiment 3: less than 1: $R^2 = .25$, $F(1, 14) = 4.609$, $p < .05$; within 0.1: $R^2 = .44$, $F(1, 14) = 10.925$, $p < .006$; within 0.01: $R^2 = .41$, $F(1, 14) = 9.582$, $p < .009$.

Including information about violations had little effect on predictive ability. When predictions were based on a representation that included information about strength of property pair correlation but excluded information about violations, predictions were reliable: less than 1: $R^2 = .32$, $F(1, 14) = 6.610$, $p < .03$; within 0.1: $R^2 = .46$, $F(1, 14) = 11.715$, $p < .005$; within 0.01: $R^2 = .40$, $F(1, 14) = 9.233$, $p < .01$. These predictions can be compared to those based on both violations and weighting, presented at the end of the previous paragraph. This difference between humans and the model may have resulted because of a reduced proportion of violations in the model concepts. According to the correlated properties representation involving all 190 normed

entities, in the 34 biological kind prime-target pairs used in the human experiment, the cases in which one concept contained a correlated property pair but the other violated it (mean = 94.0, se = 2.7) were almost three times as frequent as cases in which concepts shared a correlated pair (mean = 33.4, se = 2.1). When the same numbers were computed for the 18 biological kind prime-target pairs using the correlated properties representation constructed from the 84 concepts on which the model was trained, there were more shared correlated pairs (mean=53.6, se = 9.9) than cases where one concept contained the pair but the other violated it (mean = 39.5, se = 5.0). Finally, when human priming effects for only the 18 biological kind pairs included in the model were predicted using representations based on the 84 concepts, the best predictor excluded information about correlational violations. Thus, it seems reasonable that the influence of violations was reduced in the model.

The main difference between human and model priming effects was that model predictions did not fail when unique properties were included. The model's biological kind priming effects were best predicted when concept similarity was computed from a correlated properties representation in which all properties were used, only correlated pairs were included, pairs were weighted by correlational strength, and information about violations was excluded: less than 1: R^2 = .49, $F(1, 14) = 13.644$, $p < .003$; within 0.1: R^2 = .63, $F(1, 14) = 23.631$, $p < .001$; within 0.01: R^2 = .63, $F(1, 14) = 23.894$, $p < .001$.

In summary, as in Experiment 3, individual and correlated properties representations based on the 84 concepts on which the model was trained were used to predict model priming effects. The analyses demonstrated that the same basic factors predicted priming effects for humans and the model. Some aspects of the correlated properties representation that control predictive performance differed, but these inconsistencies are understandable when it is considered that the representations on which predictions were based contained all of the model's knowledge.

7.0 CONVERGENCE LATENCY

A number of people have claimed that, although connectionist models often adequately account for data, they rarely provide the basis for novel predictions (Massaro & Cohen, 1991; McCloskey, 1991). In this section, we outline and empirically test a nonobvious prediction derived from the model. In our modeling, we assumed that convergence latency is monotonically related to the time required for a human to compute a concept and, therefore, to respond in a semantic task. A central characteristic of the model is that the number of iterations required to compute a concept is influenced by the intercorrelational density of a concept's properties. In the model, property pair correlational strength is directly related to connection strength. Furthermore, the larger are the weights interconnecting a concept's properties, the greater is the net input to each of that concept's conceptual units on each iteration, and the faster the concept's conceptual units are pushed to their maximum activation. Thus, the more strongly a concept's properties are intercorrelated, the larger the increases in property node activation due to mutual facilitation on each update. In summary, the more densely intercorrelated are the properties of a concept, the faster that the concept converges.

To illustrate the influence of property intercorrelations on convergence latency, two groups of 30 concepts were chosen from the model. One group consisted of concepts whose properties were densely intercorrelated; the other group consisted of concepts whose properties were sparsely intercorrelated. For example, LETTUCE has properties such *is edible*, *is green*, *used for salads*, and *is crunchy* that tend to occur together. In contrast, FREEZER has properties such as *has shelves*, *has a motor*, *is cold*, and *made of metal* whose combination is idiosyncratic. The two groups were equated for the number of individual properties per concept.

Figure 7 shows mean error as a function of time step. The computation into semantic space from orthography is labeled as time step 1. Mean error for the two groups of concepts is similar at this point. However, during time steps two through five, error for concepts with densely intercorrelated properties drops more quickly; that is, these concepts were faster to converge. During time steps six through ten, mean error is again similar for the two groups. It should be

similar in this range; the claim is not that concepts with sparsely intercorrelated properties cannot be computed, it is simply that property intercorrelational density influences how quickly a concept converges. An analysis of variance was conducted using error as the dependent variable and intercorrelational density and time step as the independent variables. Collapsed over time step, error was significantly lower for concepts with densely intercorrelated properties, $F(1, 58) = 13.070$, $p < .0007$. Furthermore, group interacted with time step, $F(9, 522) = 12.408$, $p < .0001$. Simple main effects revealed that mean error was not significantly different at time step 1, $F < 1$. However, error for densely intercorrelated concepts was lower at time steps two, three, and four ($p < .001$ in all cases), and marginally lower at time step 5 ($p < .08$). No significant differences were found over the final five time steps ($F < 1$ in all cases).

Figure 7 about here

Given these results, it was predicted that subjects would be faster to make a semantic decision to entities with densely intercorrelated properties than to matching ones with weakly intercorrelated properties.

7.1 Experiment 4 - Intercorrelational Density

For the human experiment, two groups of 42 concepts were chosen that differed in terms of intercorrelational density of their properties (dense versus sparse). The groups were equated in terms of a number of variables that may influence word recognition: printed word frequency, rated familiarity of the entity, length in letters, bigram frequency, and number of individual properties. A concreteness decision task was used in which subjects responded "yes" if the word referred to something that was a concrete object in the world (e.g., LETTUCE or FREEZER), or "no" otherwise (e.g., LOVE or MOTIVATION). This task satisfied two criteria: the decision was based on conceptual knowledge, so that it was sensitive to the computation of word meaning; and, the response was "yes" for all of the entities involved (which ranged through all 10 categories). Given the manner in which the model computed word meaning from orthography, it was predicted that Ss' decision latencies would be faster for words that referred to entities with densely intercorrelated properties.

Method

Subjects. Ten University of Rochester undergraduates participated for course credit.

Materials. Stimuli consisted of 2 groups of 42 concepts. One group consisted of concepts with densely intercorrelated properties, and the other consisted of concepts with sparsely intercorrelated properties (see the Appendix for the stimuli). Intercorrelational density of a concept was computed by summing the percentage of shared variance between each of its property pairs, using only the properties that occurred in a minimum of three concepts. Mean concept intercorrelational density was higher for the dense group (mean = 938, se = 45) than for the sparse group (mean = 145, se = 15), $t(82) = 16.717$, $p < .0001$. The dense group contained 17 biological kinds and 25 artifacts; the sparse group contained 3 biological kinds and 39 artifacts. The groups were equated on 5 variables that may influence word recognition (means and standard errors in parentheses): printed word frequency (Kucera & Francis, 1967): dense (19, 7), sparse (18, 4), $t(82) = 0.114$, $p > .9$; rated familiarity of the object as judged by 20 subjects in a separate norming study: dense (5.5, 0.1), sparse (5.5, 0.1), $t(82) = -0.394$, $p > .6$; length in letters: dense (5.5, 0.3), sparse (5.5, 0.2), $t(82) = 0.0$; positional bigram frequency (Solso & Juel, 1980): dense (1016, 103), sparse (990, 99), $t(82) = 0.185$, $p > .8$; and the number of individual

properties listed by subjects in the norms: dense (16, 0.3), sparse (16, 0.4), $t(82) = 0.669$, $p > .5$.

Eighty-four abstract concepts were chosen to serve as "no" responses. They included words such as LOVE and MOTIVATION that do not refer to concrete entities. Fifteen concrete words and 15 abstract words, not included in the experimental items, were used for practice.

Procedure. A 30-item practice session was provided. Subjects were instructed to press the "yes" key (always beneath their dominant hand) as quickly and accurately as possible if the word referred to a concrete object, or press the "no" key (always beneath their nondominant hand) if it did not. There were 168 trials in the main section, 2 blocks of 84 trials separated by a short break. Order of the blocks was counterbalanced across subjects. Trials within a block were presented in random order. Each trial proceeded as follows: a 1500 ms intertrial interval; an asterisk in the center of the screen for 500 ms; a blank screen for 100 ms; a concept name until the subject responded. It took approximately 7 minutes for a subject to complete the experiment.

Design. There was one factor, property intercorrelational density, with 2 levels, dense and sparse. Intercorrelational density was a within subjects and a between items factor.

Results

Latency and accuracy of subjects' responses were recorded. All response latencies greater than 3 standard deviations from the mean of the correct test trials were replaced by the cutoff value (1.6% of the trials). One-tailed paired-samples t-tests were used to determine the effect of intercorrelational density on decision latency. Subjects were faster to judge that a word referred to a concrete object if the properties of the object were densely intercorrelated (mean = 604 ms, standard error = 7 ms) than if they were sparsely intercorrelated (mean = 629 ms, standard error = 10 ms), $t(9) = 3.45$, $p < .004$ by subjects, $t(82) = 2.04$, $p < .03$ by items.

Subjects committed errors on 3.9% of the test trials, and 6.9% of the negative filler trials. Subjects committed 3.6% errors to concepts that were densely intercorrelated and 4.3% errors to sparsely intercorrelated concepts. This difference was significant by subjects, $t(9) = 1.923$, $p < .05$, but not by items, $t(82) = 0.236$, $p > .4$.

Discussion

Convergence latency in the model predicted human decision latency in a concreteness decision task. Both measures are sensitive to encoded relationships among conceptual elements. This result provides further evidence that the computation of conceptual representations from word form depends, in part, on encoded knowledge of the higher order structure of semantic space. These results provide further support for a conceptualization of the access of word meaning in which the computation from word form to conceptual memory puts the system into the correct area of semantic space, then semantic correlational structure drives the system to a stable state. This prediction is not unique to our model; it follows from any model in which co-occurrences among semantic properties or microfeatures are encoded and directly influence convergence latency (e.g., Hinton & Shallice, 1991; Masson, 1992; Plaut, 1991; Sharkey, 1989).

8.0 GENERAL DISCUSSION

The empirical work and associated modeling reported in this article made four main contributions. First, the property verification task demonstrated that correlations among properties of real-world entities are encoded in conceptual memory. Second, property intercorrelations were shown to be integral to the contrast between artifacts and biological kinds. In an analysis of semantic production norms, properties of biological kind concepts were more densely intercorrelated than those of artifacts. Furthermore, although artifact priming effects were predicted by concept similarity in terms of individual properties, priming effects for biological kinds were predicted by correlated property pairs. Third, Experiment 3 clearly demonstrated that short interval semantic priming effects result from successively computing similar concepts. In previous priming studies, priming effects are either attributed to associative relatedness

(McNamara, 1992) or their source is not specified (see Hodgson, 1991; Neely, 1990). Experiment 3 clearly showed that short interval priming effects are predictable from property-based concept similarity. Fourth, the experiments and modeling provided evidence that encoded knowledge of the correlational structure of entities plays a key role in computing the meaning of words. The modeling illustrated that a network whose computational dynamics are dependent on attractor basins formed from encoding property intercorrelations captures some basic characteristics of human performance on semantic tasks. The results of Experiments 2 and 3 were naturally modeled using a connectionist framework, specifically, a modified Hopfield (1982, 1984) network. Also, a prediction derived from the fact that a concept's convergence latency in the model is influenced by intercorrelational density of its properties was verified in Experiment 4.

The experimental and modeling work described above is best regarded as lying at the interface between word recognition and concepts and categorization. Hence, issues were addressed that are rarely investigated in word recognition research. Most work in word recognition focuses on how an abstract representation of a word's form is contacted from auditory or written input (e.g., Becker, 1980; Marslen-Wilson & Welsh, 1978), or on how a word's phonological representation is accessed from its written form (e.g., Patterson, Marshall, & Coltheart, 1985; Seidenberg & McClelland, 1989). Clearly, computing word meaning is an important aspect of word recognition. The most theoretically important aspect of the model was the way in which basins of attraction that resulted from encoding the relationships among semantic properties were central to the computation of conceptual representations from words. This aspect is shared with other recent connectionist models of word recognition (Hinton & Shallice, 1991; Masson, 1992; Sharkey, 1989). The concept of attractor basins may be useful in understanding a number of primary word recognition phenomena. For example, the lexicon is typically conceptualized as a set of nodes, each of which corresponds to a lexical item (e.g., Morton, 1969). Alternatively, the lexicon may be conceptualized as a set of attractor basins. In this view, a lexical item corresponds to a stable state in representational space. As another example, Hinton and Shallice (1991) and Plaut (1991) have shown that attractor basins are critical to understanding phenomena that characterize deep dyslexia.

These ideas may also shed light on a recent issue of great interest to word recognition researchers, mediated priming (McNamara, 1992). It has been shown that, although two concepts are not directly associated according to free association norms (e.g., LION and STRIPES), they may prime one another (Balota & Lorch, 1986; McNamara & Altarriba, 1988). McNamara interprets these results within a semantic network theory such as Collins and Loftus (1975). He claims that LION primes STRIPES because activation of the LION node spreads to TIGER and then to STRIPES. Although the notion of association is not a primitive in the conceptual system that our model represents, it provides a plausible explanation for mediated priming effects. LION shares properties with TIGER. These properties are intercorrelated. Therefore, when LION is computed, STRIPES will be slightly activated because properties of LION are correlated with STRIPES. A small priming effect will result, as found in these studies. This explanation is currently being explored further.

Another important issue that can be explored using our model concerns the use of conceptual knowledge in the lexical decision task. Lexical decision is the most frequently used task in psycholinguistic research. However, the knowledge on which people base their lexical decisions is a matter of debate (Balota & Chumbley, 1984). Within the connectionist framework, Seidenberg and McClelland (1989) have accounted for the results of some lexical decision studies by explaining lexical decision as a familiarity decision that is based primarily on people's knowledge of how words are spelled and pronounced. However, their model predicts that people should have much greater difficulty distinguishing words from nonwords than has been found in human studies (Besner, Twilley, McCann, & Seergobin, 1990). Because the mappings within English orthography and from orthography to phonology are quasi-regular, coherent representations can be computed for novel strings if they possess subunits that capture these regularities (e.g., NUST). In contrast, because the mapping from spelling to meaning is largely

arbitrary, it is difficult to compute a meaning for a nonword (the meaning of NUST?). Thus, the mapping from orthography to conceptual memory may provide a more reliable basis for lexical decisions. This idea will be explored in future work.

8.1 Limitations

The main limitations of this work stem from the use of probabilistic models of conceptual structure. Probabilistic models are best viewed as a necessary component of a theory of concepts and categorization, but insufficient to explain the broad range of categorization behavior of which humans are capable (Medin, 1989). Although explaining the full range of human categorization behavior requires representational and processing assumptions beyond the scope of these models, they do capture a great deal of the knowledge that is applied in people's everyday use of entity concepts. Because we used production norms, our model was implemented in a way that de-emphasized the role of specific exemplars in conceptual representation. As in other models of this type (e.g., Hintzman, 1986), it was assumed that when a word such as DOG is read or heard in isolation, people compute conceptual representations that contain the context-independent properties associated with the set of entities to which the word refers (Barsalou, 1982). The implemented model can be regarded as a snapshot of those representations that included knowledge of how properties co-occur.

There are a number of issues concerning which property intercorrelations are encoded that our research did not address. Our model incorporated the assumption that the mind is structured so that any pair of properties are encoded if they occur together and are simultaneously attended. Thus, the encoding of property intercorrelations may be an unconscious and automatic consequence of attending to the world around you. This view is also supported by Younger and Cohen's (1983) experiment with 10 month old infants. It is also probable, however, that people's theories and knowledge of the world partly determine what pairs of properties are noticed. In this view, people consciously notice that two properties tend to occur together in the world and then encode that fact. Barrett et al. (1992) have shown that a child's sensitivity to correlated properties was modulated by whether the relationship made sense. We believe that people automatically encode property intercorrelations as a natural by-product of attending to entities, but that more complex types of causal knowledge can result in certain relationships being emphasized. It is also possible that automatic encoding of correlations gives rise to theories of the world. For example, a child may notice that things that fly tend to have wings. This may lead the child to hypothesize a causal relationship, resulting in a question to a parent concerning the function of wings. In any case, it will be important to better understand the relationship between conscious and unconscious aspects of the encoding of correlations/relationships among properties.

In closing, a limitation of this and other work in categorization is its reliance on simple, concrete entities. Because they are particularly tractable for study, concrete entities have been used as stimuli in virtually every study of real-world concepts. It is not clear whether the conclusions from these studies will generalize to more abstract types of concepts such as emotions, events, or actions. However, the study of real-world entity concepts is important and interesting in its own right.

REFERENCES

Anderson, J. R. (1973). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, **98**, 409-429.

Ashcraft, M. H. (1978). Property dominance and typicality effects in property statement verification. *Journal of Verbal Learning and Verbal Behavior*, **17**, 155-164.

Balota, D. A., & Chumbley, J. (1984). Are lexical decisions a good measure of lexical access? The role of word-frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 340-357.

Balota, D. A, & Lorch, R. F. (1986). Depth of automatic spreading activation: Mediated priming effects in pronunciation but not lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12**, 336-345.

Barrett, S. E., Abdi, H., Murphy, G. L., & Gallagher, J. M. (in press). Theory-based correlations and their role in children's concepts. *Child Development*.

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

Battig, W. F., & Montague, W. E. (1969). Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. *Journal of Experimental Psychology Monograph*, **80**, 1-46.

Becker, C. A. (1980). Semantic context effects in visual word recognition: An analysis of semantic strategies. *Memory & Cognition*, **8**, 493-512.

Beiderman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.

Besner, D., Twilley, L., McCann, R. S., & Seergobin, K. (1990). On the association between connectionism and data: Are a few words necessary? *Psychological Review*, **97**, 432-446.

Boyd, R. (1984). Natural kinds, homeostasis, and the limits of essentialism. Paper presented at Cornell University.

Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, **13**, 475-511.

Clark, H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, **12**, 335-359.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, **82**, 407-428.

de Groot, A. M. B. (1984). Primed lexical decision: Combined effects of the proportion of related prime-target pairs and the stimulus-onset asynchrony of prime and target. *The Quarterly Journal of Experimental Psychology*, **36A**, 253-280.

Dell, G. S. (1986). A spreading-activation theory of retrieval in language production. *Psychological Review*, **93**, 283-321.

den Heyer, K., Briand, K., & Dannenbring, G. L. (1983). Strategic factors in a lexical-decision task: Evidence for automatic and attention-driven processes. *Memory & Cognition*, **11**, 374-381.

Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, **120**, 339-357.

Favreau, M., & Segalowitz, N. S. (1983). Automatic and controlled processes in first- and second-language reading of fluent bilinguals. *Memory & Cognition*, **11**, 565-574.

Fischler, I. (1977). Semantic facilitation without association in a lexical decision task. *Memory & Cognition*, **5**, 335-339.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and animate-inanimate distinctions as examples. *Cognitive Science*, **14**, 79-106.

Gelman, S. A., & Kremer, K. E. (1991). Understanding natural cause: Children's explanations of how objects and their properties originate. *Child Development*, **62**, 396-414.
Gluck, M. A. (1991). Stimulus generalization and representation in adaptive network models of category learning. *Psychological Science*, **2**, 50-55.
Gluck, M. A., & Bower, G. H. (1988). Evaluating an adaptive network model of human learning. *Journal of Memory and Language*, **7**, 166-195.
Gluck, M. A., Bower, G. H., & Hee, M. R. (1989). A configural-cue network model of animal and human associative learning. In *Proceedings of the Eleventh annual Conference of the Cognitive Science Society*, Ann Arbor, Michigan. August 16-19, 1989 (pp. 323-332). Hillsdale, NJ: Erlbaum.
Gough, P. B., & Cosky, M. J. (1977). One second of reading again. In N. J. Catstellan, Jr., D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive Theory* (Vol. 2). Hillsdale, NJ: Erlbaum.
Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, **18**, 441-461.
Hayes-Roth, B. & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning and Verbal Behavior*, **16**, 321-338.
Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.
Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation. Santa Fe Institute Studies in the Sciences of Complexity Lecture Notes (Vol. 1)*. Workingham, England: Addison-Wesley.
Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, **98**, 74-95.
Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, **93**, 411-428.
Hintzman, D. L., & Ludlam, G. (1980). Differential forgetting of prototypes and old instances: Simulation by an exemplar-based classification model. *Memory & Cognition*, **8**, 378-382.
Hodgson, J. M. (1991). Informational constraints on pre-lexical priming. *Language and Cognitive Processes*, **6**, 169-264.
Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Science*, **79**, 2254-2558.
Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Science*, **81**, 3088-3092.
Jared, D., & Seidenberg, M. S. (1992). Does word identification proceed from spelling to sound to meaning? *Journal of Experimental Psychology: General*, **120**, 358-394.
Keil, F. C. (1987). Conceptual development and category structure. In U. Neisser (Ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization* (pp. 175-200). Cambridge, MA: Cambridge University Press.
Keil, F. C. (1989). *Concepts, Kinds, and Cognitive Development.* Cambridge, MA: MIT Press.
Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22-44.
Kucera, H., & Francis, W. N. (1967). *Computational Analysis of Present-day American English*. Providence, RI: Brown University Press.
Lewis, J. L. (1970). Semantic processing of unattended messages using dichotic listening. *Journal of Experimental Psychology*, **85**, 225-228.
Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior*, **23**, 250-269.
Marslen-Wilson, W. D., & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, **10**, 29-63.
Massaro, D. W., & Cohen, M. M. (1991). Integration versus interactive activation: The joint influence of stimulus and context in perception. *Cognitive Psychology*, **23**, 558-614.

Masson, M. E. J. (1992). A distributed memory model of semantic priming. manuscript submitted for publication.
McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, **88**, 375-407.
McClelland, J. L., & Rumelhart, D. E. (1986). *Parallel Distributed Processing. Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT press.
McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. *Psychological Science*, **2**, 387-395.
McCloskey, M. E., & Glucksberg, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, **6**, 462-472.
McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **18**, 1155-1172.
McNamara, T. P. (1992). Priming and constraints it places on theories of memory and retrieval. *Psychological Review*, **99**, 650-662.
McNamara, T. P., & Altarriba, J. (1988). Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *Journal of Memory and Language*, **27**, 545-559.
McNamara, T. P., & Sternberg, R. J. (1983). Mental models of word meaning. *Journal of Verbal Learning and Verbal Behavior*, **22**, 449-474.
McRae, K. (1991). *Independent and correlated properties in artifact and natural kind concepts*. Unpublished doctoral dissertation, McGill University, Montreal.
Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, **44**, 1469-1481.
Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **8**, 37-50.
Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.
Medin, D. L., & Schwanenflugel, P. J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Learning, Memory, and Memory*, **7**, 355-368.
Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, **76**, 165-178.
Murdock, B. B. (1987). Serial-order effects in a distributed-memory model. In D. S. Gorfein & R. R. Hoffman (Eds.), *Memory and learning: The Ebbinghaus Centennial Conference* (pp. 227-310). Hillsdale, NJ: Erlbaum.
Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, **92**, 289-316.
Murphy, G. L., & Wisniewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in Cognitive Science* (Vol. 2): *Theory and Applications* (p. 23-45). Chichester, England: Ellis Horwood.
Myers, J. L. (1979). *Fundamentals of experimental design*. Boston, MA: Allyn and Bacon.
Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, **106**, 22-254.
Neely, J. H. (1990). Semantic priming effects in visual word recognition: A selective view of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic Processes in Reading: Visual Word Recognition* (pp. 264-336). Hillsdale, NJ: Lawrence Erlbaum.
Neely, J. H., & Keefe, D. E. (1989). Semantic context effects on visual word processing: A hybrid prospective/retrospective processing theory. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. **24**). New York: Academic Press.

Neuman, P. G. (1974). An attribute frequency model for the abstraction of prototypes. *Memory & Cognition*, **2**, 241-248.
Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **10**, 104-114.
Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**, 39-57.
Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch and B. B. Lloyd (Eds.). *Cognition and categorization* (pp. 259-303). Hillsdale, NJ: Lawrence Erlbaum.
Patterson, K. E., Marshall, J. C., & Coltheart, M. (1985). *Surface Dyslexia: Neuropsychological and Cognitive Studies of Phonological Reading*. London: Erlbaum.
Plaut, D. C. (1991). *Connectionist neuropsychology: The breakdown and recovery of behavior in lesioned attractor networks*. Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh.
Posner, M. I., & Snyder, C. R. R. (1975). Facilitation and inhibition in the processing of signals. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V*. New York: Academic Press.
Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, **95**, 385-408.
Rayner, K. (1978). Eye movements in reading and information processing. *Psychological Bulletin*, **85**, 618-660.
Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 21-59). Cambridge, MA: Cambridge University Press.
Rips, L. J. , Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, **12**, 1-21.
Rolls, E. T. (1989). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin, C. Miall, & F. Mitchison (Eds.), *The Computing Neuron* (pp. 125-159). Workingham, England: Addison-Wesley.
Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, **104**, 192-233.
Rosch, E. (1978). Principles of categorization. In E. Rosch and B. B. Lloyd (Eds.). *Cognition and Categorization* (pp. 27-48). Hillsdale, NJ: Lawrence Erlbaum.
Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, **7**, 573-605.
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, **8**, 382-439.
Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, **96**, 523-568.
Sharkey, N. E. (1989). The lexical distance model and word priming. In *Proceedings of the Eleventh annual Conference of the Cognitive Science Society*, Ann Arbor, Michigan. August 16-19, 1989 (pp. 860-867). Hillsdale, NJ: Erlbaum.
Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts*. Cambridge, MA: Harvard University Press.
Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, **12**, 485-527.
Smith, L. C., Briand, K., Klein, R. M., & den Heyer, K. (1987). On the generality of Becker's verification model. *Canadian Journal of Psychology*, 41, 379-386.
Solso, R. L., & Juel, C. L. (1980). Positional frequency and versatility of bigrams for two- through nine-letter English words. *Behavior Research Methods and Instrumentation*, **12**, 297-343.

Tsodyks, M. V., & Feigelman, M. V. (1988). The enhanced storage capacity in neural networks with low activity level. *Europhysics Letters*, **6**, 101-105.
Tversky, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
Warren, R. E. (1977). Time and the spread of activation in memory. *Journal of Experimental Psychology: Human Learning and Memory*, **3**, 458-466.
Younger, B. A., & Cohen, L. B. (1983). Infant perception of correlations among attributes. *Child Development*, **54**, 858-867.

APPENDIX

Experiment 1 and 3 stimuli: The complete set of stimuli in pairs of similar items.

BIRD	robin	eagle	hawk
parakeet	budgie	vulture	buzzard
finch	canary	stork	crane
duck	chicken	starling	crow
goose	turkey	ostrich	emu
MAMMAL	pig	tiger	lion
cat	dog	goat	sheep
mouse	rat	horse	pony
caribou	moose	whale	dolphin
deer	fawn	bull	cow
FRUIT	apple	lime	lemon
peach	nectarine	honeydew	cantaloupe
plum	prune	pineapple	coconut
grapefruit	orange	cherry	cranberry
grape	raisin	tangerine	mandarin
VEGETABLE	corn	onions	garlic
cauliflower	broccoli	yam	potato
carrot	celery	squash	pumpkin
lettuce	cabbage	peas	beans
cucumber	zucchini	radish	beets
CLOTHING	socks	nylons	leotards
dress	skirt	mittens	gloves
blouse	shirt	trousers	pants
bra	camisole	boots	shoes
tie	belt	sandals	slippers
FURNITURE	chair	dresser	closet
couch	sofa	cupboard	shelves
bureau	desk	cushion	pillow
lamp	chandelier	curtains	drapes
radio	stereo	carpet	mat
KITCHEN	stove	freezer	fridge
plate	saucer	microwave	toaster
mixer	blender	pan	pot
fork	spoon	faucet	sink
mug	cup	jar	bottle
TOOL	hammer	pencil	crayon
drill	screwdriver	hoe	shovel
file	sandpaper	screws	bolts
wrench	pliers	level	ruler
vice	clamp	shed	barn

VEHICLE	car	jet	airplane
subway	bus	canoe	raft
wagon	cart	yacht	ship
truck	van	jeep	dunebuggy
tricycle	bike	scooter	motorcycle
WEAPON	gun	slingshot	catapult
dagger	knife	stone	rock
sword	spear	axe	tomahawk
pistol	rifle	cannon	bazooka
bomb	missile	stick	club

Experiment 2 stimuli: Strong and weakly intercorrelated groups.

Property	**Strong**	**Weak**
come (comes) in different colours	pants	carpet
eats	parakeet	whale
found in bathrooms	faucet	mat
found in kitchens	microwave	faucet
has an engine	motorcycle	yacht
has a handle	tomahawk	wagon
has a long tail	mouse	pony
has a seat	tricycle	chair
has leaves	lettuce	pineapple
has legs	caribou	canary
has skin	apple	potato
has teeth	lion	rat
has wheels	scooter	cannon
hunted by people	deer	ducks
is (are) comfortable	couch	slippers
is (are) decorative	drapes	carpet
is (are) transparent	jar	nylons
is breakable	plate	crayon
is colourful	budgie	carpet
is crunchy	carrot	apple
is dangerous	rifle	motorcycle
is electrical	blender	drill
is fast	car	microwave
is green	lettuce	lime
is round	peach	cabbage
is soft	sofa	carpet
is sweet	peach	raisin
is tropical	coconut	parakeet
made of material	pants	couch
made of plastic	jar	sandals
migrates	goose	caribou
used as a pet	parakeet	dog
used as a weapon	axe	stick
used for protection	pistol	dog
used for transportation	van	raft
used in the circus	lion	cannon
worn for warmth	gloves	shirt

Experiment 4 stimuli: Densely and Sparsely intercorrelated concepts.

Dense	**Dense**	**Sparse**	**Sparse**
apple	gun	airplane	missile
axe	hawk	barn	nylons
blouse	jar	belt	onions
boots	lemon	bike	pencil
bottle	lettuce	blender	radio
broccoli	moose	bolts	raft
cabbage	pants	cannon	raisin
canary	parakeet	canoe	sandals
cantaloupe	peach	carpet	sandpaper
car	pineapple	chandelier	screws
carrot	plate	curtains	ship
cauliflower	plum	cushion	shoes
celery	radish	desk	shovel
chair	robin	drapes	sink
cherry	shirt	drill	skirt
crow	socks	faucet	spoon
cucumber	sofa	fork	stick
dagger	spear	garlic	subway
eagle	stork	hoe	tie
goose	sword	jeep	wrench
grape	van	lamp	yacht

Figure Captions

Figure 1. Schematic illustration of concepts in terms of individual and correlated properties. The large circles, labeled A and B, represent concepts and the small circles represent properties. Connected circles represent correlated property pairs.

Figure 2. Example of an item for Experiment 2. The target property, *hunted by people*, was correlated with 11 properties of DEER (strong group), but only 4 properties of DUCK (weak).

Figure 3. Basic architecture of the model. The conceptual units were fully bidirectionally interconnected (not all connections are shown). The orthographic units were unidirectionally connected to the conceptual units and were not interconnected.

Figure 4. Convergence patterns for the average of the 84 concepts, as well as CRAYON, the most slowly converging concept, and JET, which converged then activated additional properties.

Figure 5. Mean activation of target properties for the strong and weakly intercorrelated groups. Properties that are more strongly intercorrelated with others of a concept received a boost in activation.

Figure 6. After clamping the concept then activating the target property, the computed pattern drifted further away from the concepts of the weakly intercorrelated group, as measured by error (6a) and number of correctly activated units (6b).

Figure 7. Convergence patterns for concepts with densely and sparsely intercorrelated properties. At time steps 2 through 5, error is less for densely intercorrelated properties.

Figure 1

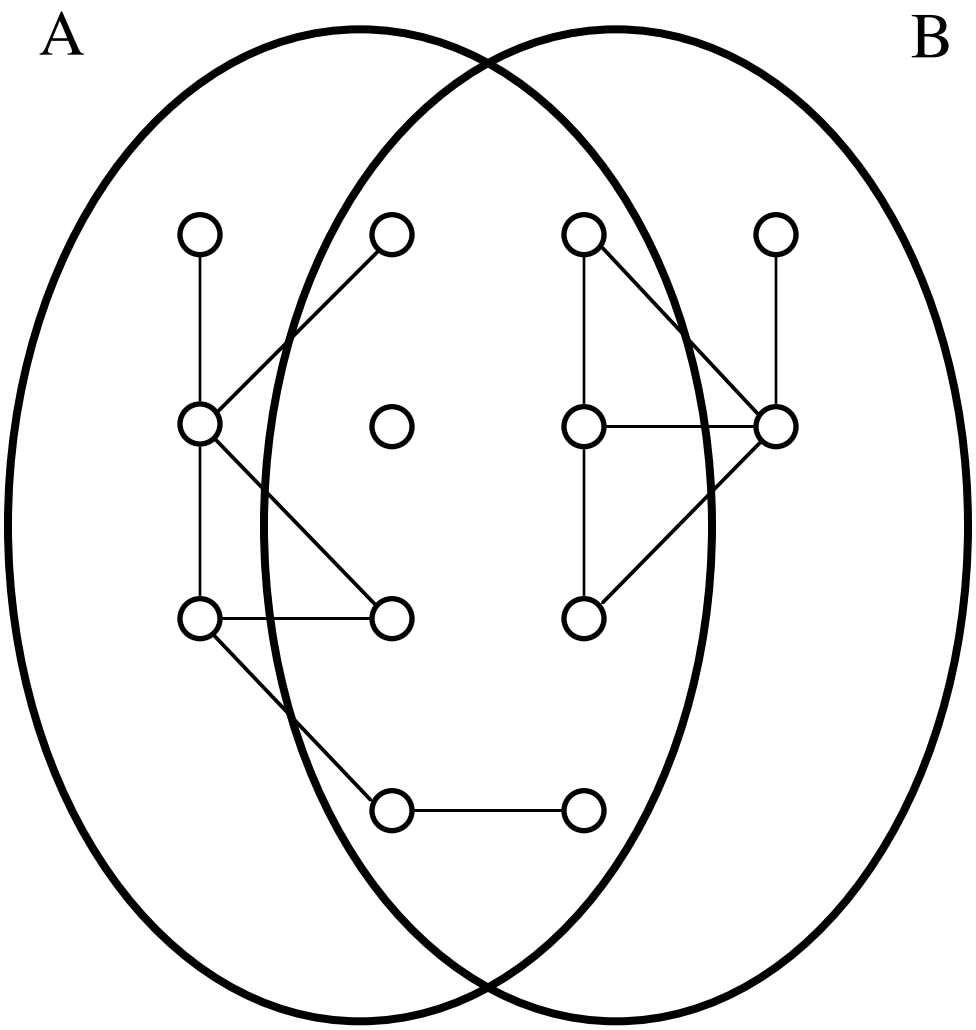

Figure 2

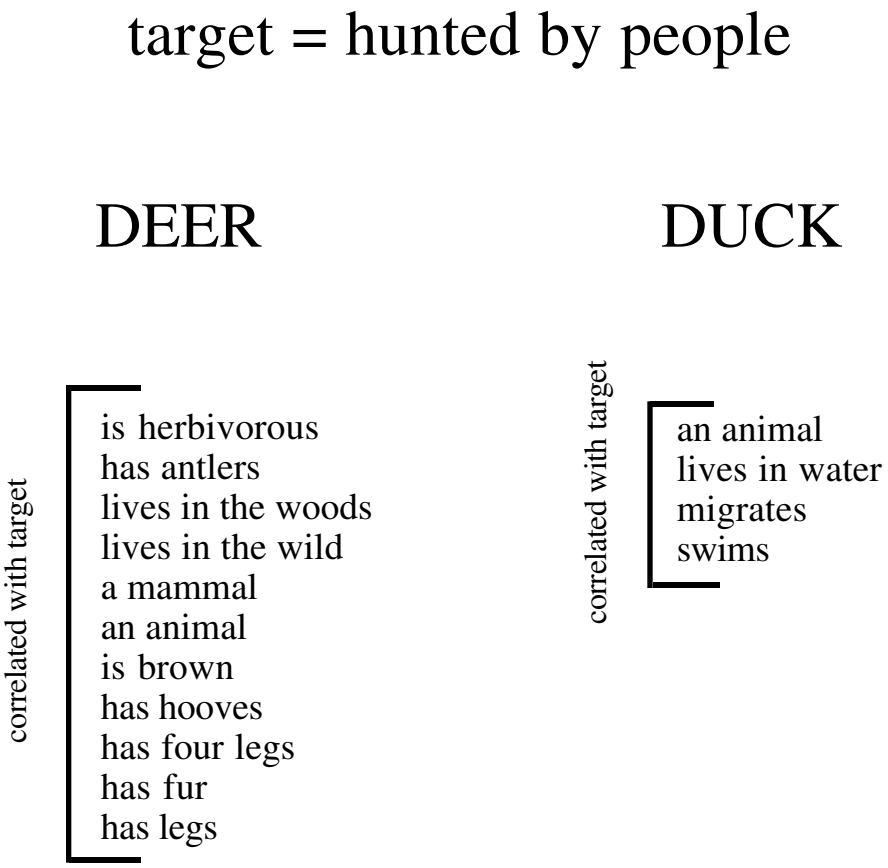
target = hunted by people
DEER
DUCK
correlated with target
is herbivorous
has antlers
lives in the woods
lives in the wild
a mammal
an animal
is brown
has hooves
has four legs
has fur
has legs
correlated with target
an animal
lives in water
migrates
swims

Figure 3

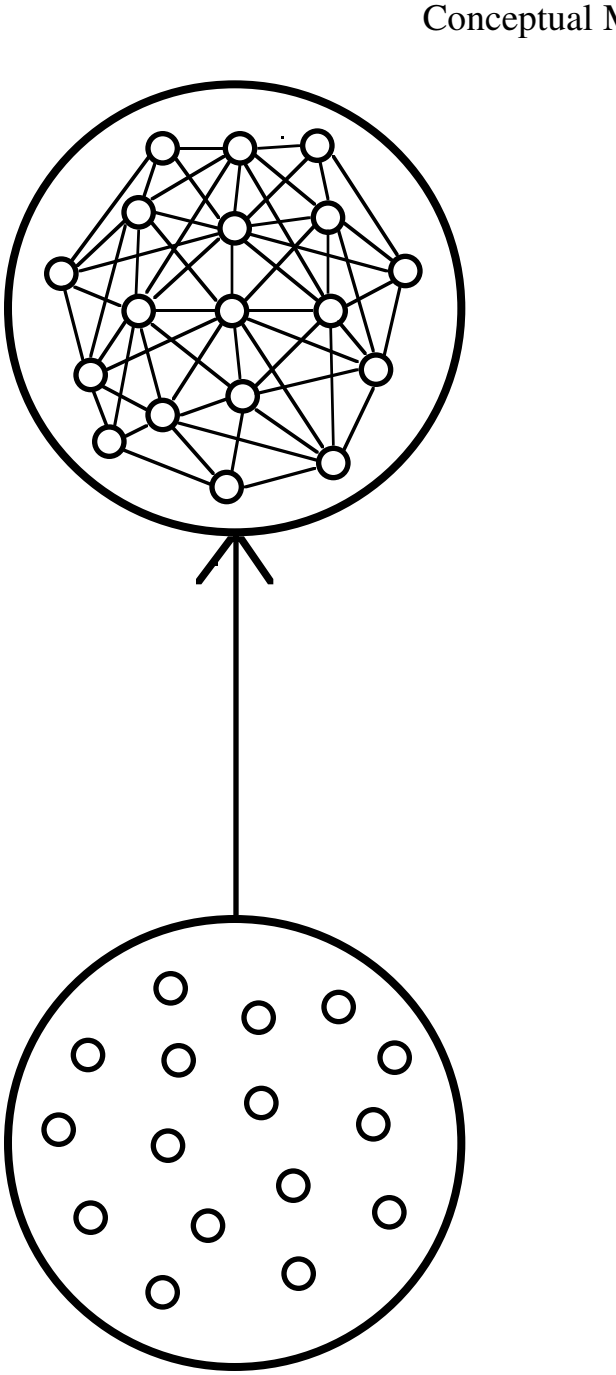

Figure 4

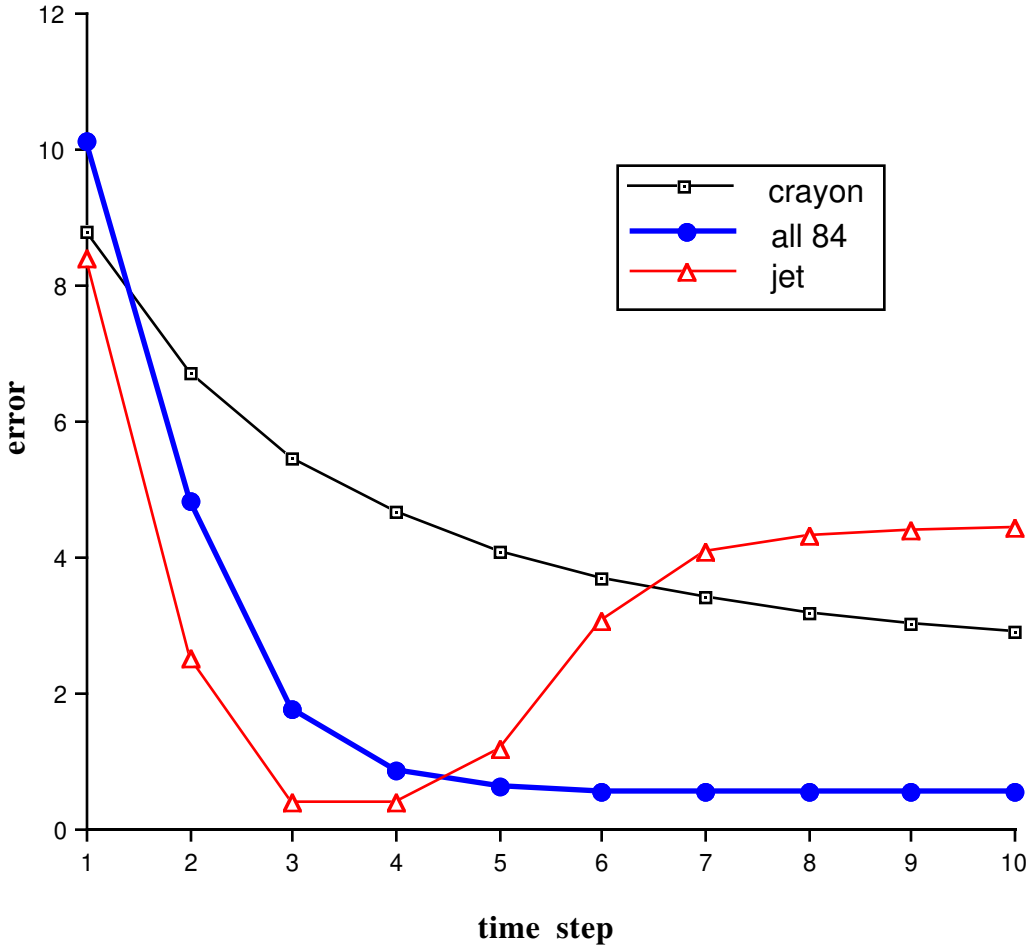

Figure 5

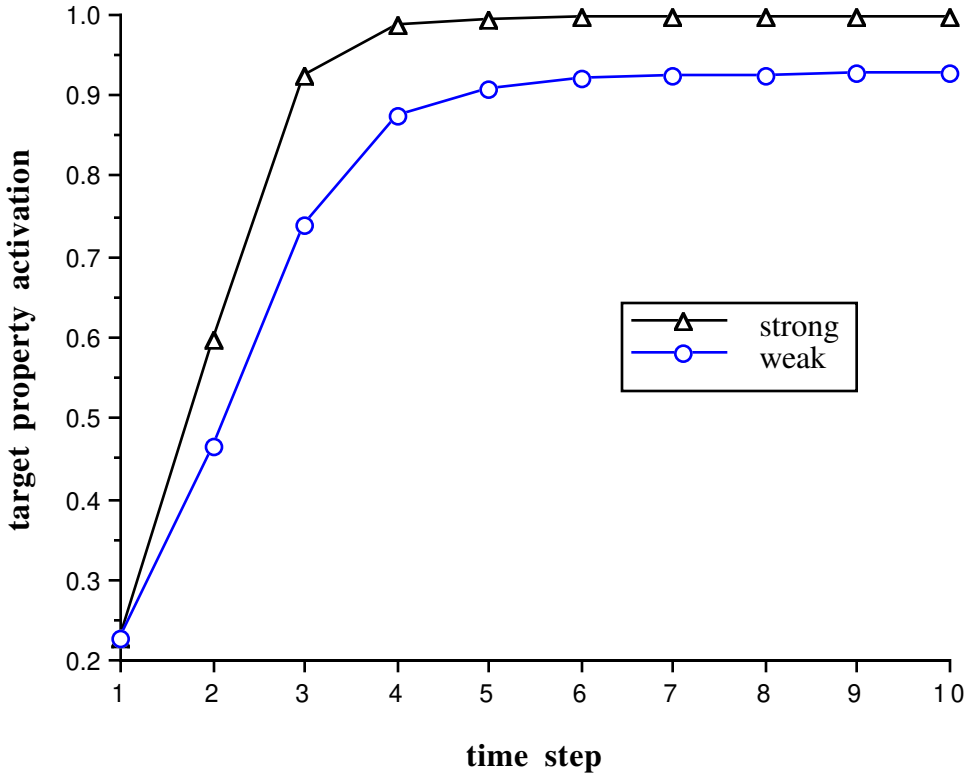
target property activation
time step
strong
weak
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
1
2
3
4
5
6
7
8
9
10

Figure 6a

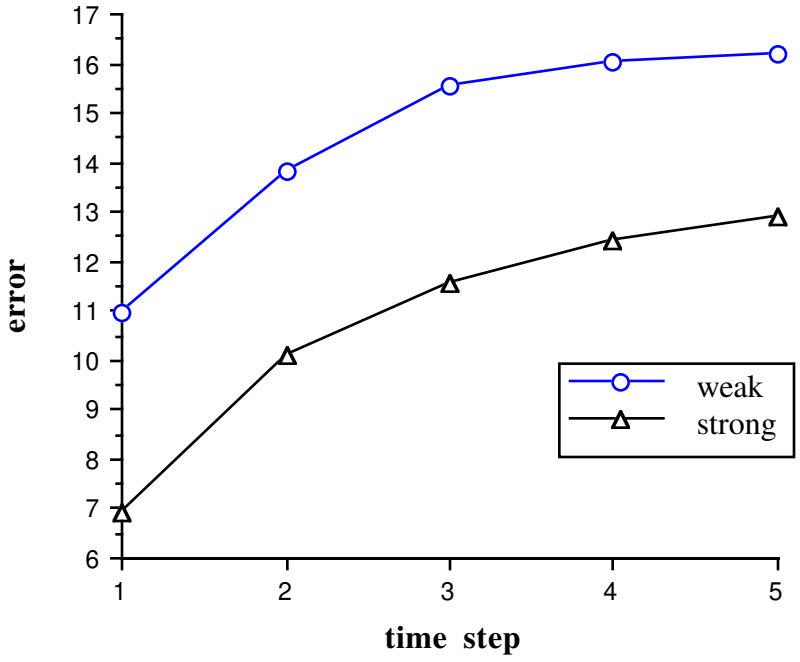

Figure 6b

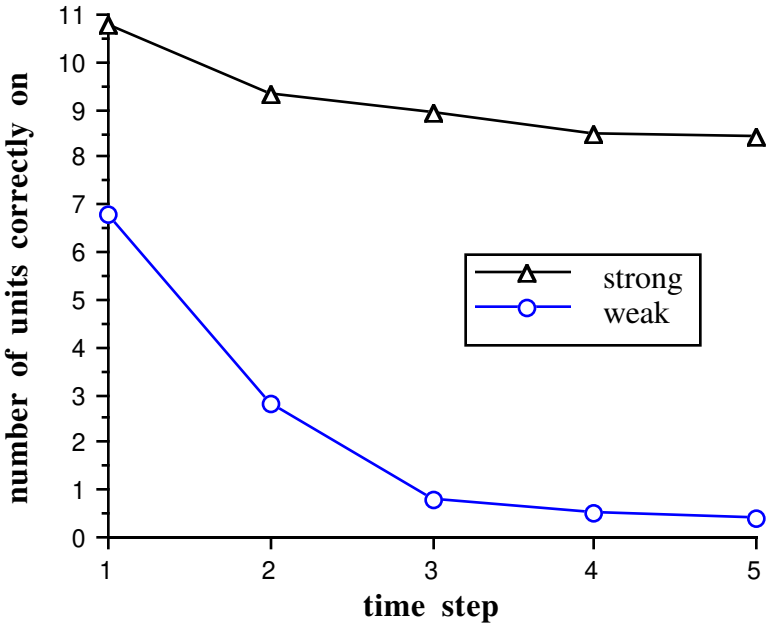

Figure 7

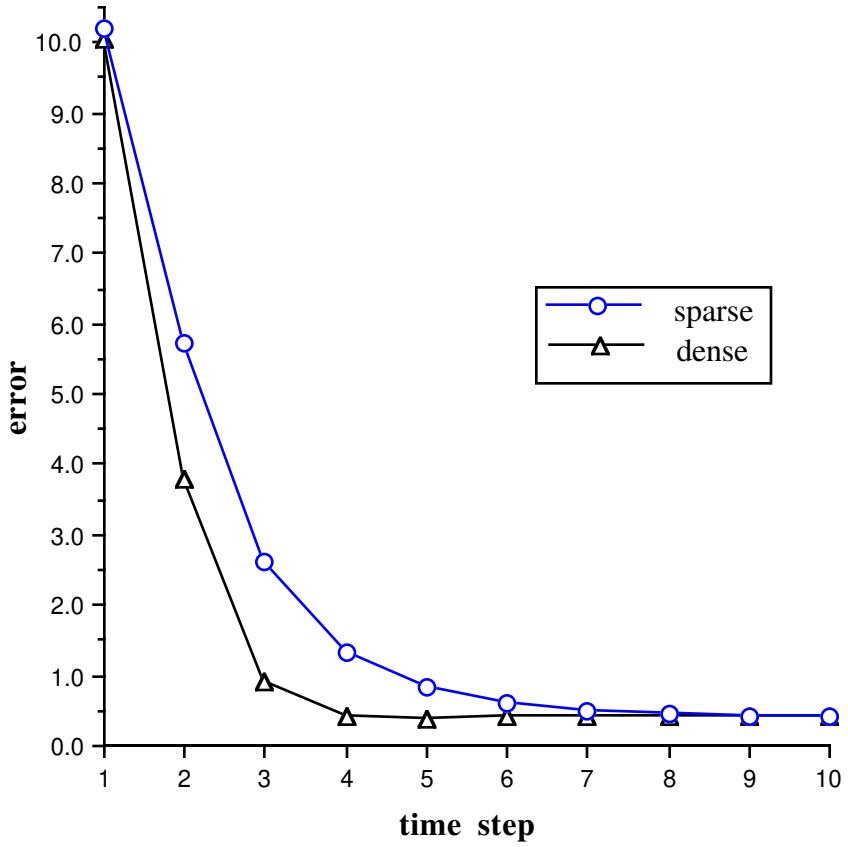
10.0
9.0
8.0
7.0
6.0
5.0
4.0
3.0
2.0
1.0
0.0
error
1
2
3
4
5
6
7
8
9
10
time step
sparse
dense

Table 1
Design of filler trials for Experiment 3

	Similar Trials		Dissimilar Trials	
	Target			
Prime	animate	inanimate	animate	inanimate
animate	**caribou** **moose**		pilot salamander	camel magazine frog tunnel
inanimate		square triangle	**drill** **moose** paper astronaut	pen gravel

Note: Each prime-target pair corresponded to 20 trials in the "is it an object?" task, 10 in all others. Items in bold were included in analyses, all others were fillers.

Table 2
Mean decision latencies and priming effects (in ms) by task for Experiment 3

Task	Prime Type	Mean[a]	Priming Effect
Animate	Similar	692 (17)	63
	Dissimilar	755 (21)	
Object	Similar	733 (11)	50
	Dissimilar	783 (12)	
Grow	Similar	741 (13)	23
	Dissimilar	764 (15)	
Man-made	Similar	899 (23)	53
	Dissimilar	952 (24)	

[a]standard error in parentheses