Further Evidence for Feature Correlations in Semantic Memory

KEN MCRAE, GEORGE S. CREE, and ROBYN WESTMACOTT,
University of Western Ontario
VIRGINIA R. DE SA, *University of California, San Francisco*

Abstract The role of feature correlations in semantic memory is a central issue in conceptual representation. In two versions of the feature verification task, participants were faster to verify that a feature (<is juicy>) is part of a concept (*grapefruit*) if it is strongly rather than weakly intercorrelated with the other features of that concept. Contrasting interactions between feature correlations and SOA were found when the concept versus the feature was presented first. An attractor network model of word meaning that naturally learns and uses feature correlations predicted those interactions. This research provides further evidence that semantic memory includes implicitly learned statistical knowledge of feature relationships, in contrast to theories such as spreading activation networks, in which feature correlations play no role.

Our environment is highly structured. In the domain of language processing, for instance, there are numerous sources of structure to which people are sensitive. Some words occur together more often than chance within sentences, as do some letters and phonemes within words. In this article, we focus on the fact that some semantic features tend to co-occur across objects and entities. For example, things in the world that <have wings> also tend to <have beaks> and <have feathers>.[1] Almost 25 years ago, Rosch and her colleagues noted that this environmental structure could be conceptualized in terms of co-occurring clusters of features (Rosch, 1978; Rosch & Mervis, 1975), an observation upon which the present research builds.

Although it is uncontroversial that people learn and use statistical relationships between a concept and its features (e.g., between *robin* and <has wings>; Smith & Medin, 1981), many researchers claim that semantic memory does not include knowledge of statistically based feature correlations (Murphy & Wisniewski, 1989). In contrast, some research on incidental concept learning (Billman & Knutson, 1996) and the computation of word meaning (McRae, de Sa, & Seidenberg, 1997, hereafter MdSS) has found that people do indeed learn feature correlations. The primary goal of this article is to add to this debate by presenting further evidence for the role of feature correlations in lexically based semantic tasks. The secondary goal is to use an attractor network to elucidate the principles that we feel are important for understanding the role of feature correlations in lexical processing, namely, implicit and incremental learning through experience with the environment, in conjunction with distributed semantic representations.

ARGUMENTS AGAINST FEATURE CORRELATIONS

We begin by noting that in this article, a "feature correlation" refers to a pair of features that tend to appear together in basic-level concepts. For example, the features <has feathers> and <has a beak> are correlated because they co-occur in things like robins, sparrows, and hawks. The theoretical arguments that have led researchers to claim that people do not learn these correlations are based primarily on spreading activation networks and prototype models. From the standpoint of a spreading activation network, encoding feature correlations requires linking every feature with every other feature, weighted by the degree of correlation (Smith & Medin, 1981). Similarly, if a prototype is taken as a list of features, feature correlations could be instantiated as direct links between them. However, the process of linking features is considered problematic. Because there are a huge number of possible feature pairs to consider in the world, the task of constructing links between each correlated pair is viewed as computationally intractable. Some researchers have dealt with this problem by claiming that such a link is constructed only if the correlation is explicitly noticed (Murphy & Wisniewski, 1989), and that a relationship between two features would be noticed only if a person previously possessed an underlying theory for why the features might co-occur, for example, that wings are used for flying (Murphy & Medin, 1985). This suggests that the encoding of a feature correlation is an explicit and

[1] Throughout the article, concept names and examples of stimuli are *italicized*, whereas feature names are presented in <angled brackets>.

rare event.

Empirically, two studies have failed to find robust effects of feature correlations (Malt & Smith, 1984; Murphy & Wisniewski, 1989). Using an intentional category learning task, Murphy and Wisniewski found no evidence that participants based categorization decisions or typicality judgements on statistical relations among features that they would not have expected a priori to be correlated (e.g., <blue> and <machine washable>). They concluded that "the pre-existing connection of features appears to be a necessary requirement for noticing correlations (with these types of categories and procedures, at least)" (Murphy & Wisniewski, 1989, p. 40). These arguments and empirical results have led to the claim that "interproperty relationships are outside the boundary conditions of almost all current categorization models (including prototype, exemplar, and rule-based models). Therefore, these models currently have limited generality, and this limitation is most evident where one might most want to generalize — meaningful stimuli" (Medin & Coley, 1998, p. 417).

INCIDENTAL LEARNING OF FEATURE CORRELATIONS

The null effects of feature correlations appear to be caused by the confluence of a number of factors, including the amount of experience a person has with exemplars, the amount of structure in the domain being learned, and whether learning occurs incidentally or intentionally. In category learning experiments that have found null effects such as Murphy and Wisniewski (1989), participants were given relatively little experience with the exemplars from which this information had to be extracted. In contrast, the present experiments tapped knowledge that accrued over approximately 20 years of experience with objects and entities from various basic-level categories. Second, experiments showing null effects have used impoverished stimuli both in terms of the number of training exemplars and their complexity, whereas natural concept learning takes place in a complex world with rich structure. The empirical and modeling work of Billman and her colleagues has shown that complexity assists rather than impairs learning, in that a correlation between two features is easier to learn if it is part of a coherent system of correlations (Billman, 1989; Billman & Heit, 1988; Billman & Knutson, 1996). Finally, effects of feature correlations tend to show themselves in concept learning tasks that promote incidental rather than intentional learning (Wattenmaker, 1991). In a typical intentional learning categorization experiment, participants are presented with a series of exemplars, one at a time, and are asked to place each in one of two categories. Participants receive immediate feedback (even on the first trial), with this cycle continuing until performance reaches a predetermined criterion. However, these experiments bear little resemblance to natural concept learning, which is better characterized as incidental or observational learning. We learn by observing things, using them, and interacting with them in varied contexts. These experiences provide rich feedback that focuses the learner on more diverse and integrative aspects of the stimuli than does a task in which someone is asked to find the cue that enables distinguishing between two categories.

A number of incidental learning studies are consistent with this claim (Billman & Knutson, 1996; Clapper & Bower, 1991; Kaiser & Proffitt, 1984; Reber, 1989; Wattenmaker, 1993). For example, experiments using habituation paradigms have demonstrated that 10-month-old infants are sensitive to correlations among the visual features of objects and pictures, and that these correlations appear to play an important role in categorization and concept learning (Younger, 1990; Younger & Cohen, 1983, 1986). Thus, it may be the case that the environment's structure is captured by an incremental, implicit learning mechanism (Holyoak & Spellman, 1993).

Attractor networks are useful as examples of learning models that encode and use this type of covariation information, and can be applied to semantic processing. Attractor networks can be viewed as instantiations and extensions of prototype models. These networks are interactive, parallel processing models in which distributed representations are used for constructing stable states in a multidimensional state space. Word meaning can be represented as patterns of activation across a set of semantic feature units, so that concepts are not accessed directly from separate memory locations as discrete, local units, but are computed online as unique patterns of activation. Such models are particularly well-suited for learning the structure in a set of training patterns. For example, pairs of features that co-occur in concepts on which a model is trained will have the weight between their units strengthened, and this will influence the trajectory that the model follows through state space as it settles to an attractor state representing the meaning of a word (MdSS).

ONLINE LEXICAL PROCESSING

MdSS conducted the first study of the role of feature correlations in computing word meaning. They used semantic feature production norms to construct representations for 190 living and nonliving thing concepts in terms of individual and correlated features. MdSS showed that priming effects for pairs of living things (e.g., *eagle — hawk*) were predicted by similarity in terms of correlated feature pairs, but not in terms of individual features. In contrast, priming effects for nonliving thing pairs (e.g., *truck — van*) were predicted by similarity in terms of individual features, whereas correlated feature pairs did not predict residual variation. This difference was attributed to the fact that nonliving thing concepts possess, on average, relatively fewer correlated features than do living things (Keil, 1989), thus providing less opportunity to observe their influence. MdSS also conducted a feature verification task ("Is this feature reasonably true of this concept?") in which they found that the degree to which a specific feature was correlated with the other features of a concept was the best predictor of verification latencies. A Hopfield (1982, 1984) attractor network provided mechanistic accounts of both experiments.

The purpose of the present research is to provide further evidence for the role of feature correlations in the computation of word meaning. Experiment 1 is an extension of the MdSS feature verification task that includes two important changes. First, it involves more thorough equating of possible confounding variables. Second, an SOA manipulation is used to test the influence of feature correlations over an extended time-course. In Experiment 2, the feature name was presented prior to the concept name, thus demanding a different view of the underlying computations, and leading to different model predictions and human results. In fact, the simulations predict contrasting interactions between SOA and the influence of feature correlations, and the human data bear these out.

The Model. In this section, we describe the essentials of the model used to derive predictions for the experiments. Note that we used the identical model as in MdSS, so the full set of details can be found in their Appendix B. Figure 1 shows the model's architecture, which consists of word-form representations of basic-level concepts as input, and their semantic feature representations as output. There were 84 concepts: 10 birds, 10 mammals, 8 fruits, 10 vegetables, 8 articles of clothing, 6 types of furniture, 8 kitchen items, 8 tools, 8 vehicles, and 8 weapons. Word form was represented by including one unit for each of the 379 letter triples that occurred in those words, with spaces at the beginning and end of a word treated as characters. The resulting sparse distributed representation roughly preserved orthographic similarity and did not introduce artificial systematicity into the form to meaning mapping. Each output unit corresponded to a binary semantic feature (0 = absent, 1 = present). Thus, each concept was a distributed pattern of activation across the feature units, and all features of a concept were equally salient. In addition, all words were trained with equal frequency. The semantic representations were derived from the feature norms of MdSS Experiment 1. Only 84 of the 190 normed concepts were included in the model because Hopfield networks have limited storage capacity due to the simplicity of the learning rule (Hertz, Krogh, & Palmer, 1991; Hopfield, 1982, 1984).

The network was trained to settle on a word's semantic representation given its word form. The word-form units were fully unidirectionally connected to the semantic units, but were not interconnected. The semantic units were fully bidirectionally interconnected. For the present purposes, it is important to note that because the weights were learned via the Hopfield (1982, 1984) learning rule (slightly modified

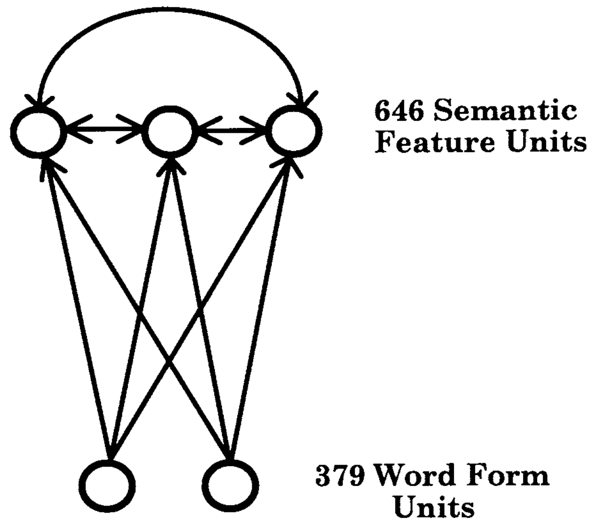

Figure 1. Architecture of the attractor network. Not all units are shown.

for sparse patterns, see MdSS), a weight connecting a pair of feature units essentially encodes the correlation between them. Thus, the model embodied two principles key to the present article: it naturally learned how features co-occur in the concepts on which it was trained, and it used this knowledge to drive the system to a stable state, that is, to compute word meaning from word form.

Experiment 1: Concept-Feature Verification

Investigating the influence of feature correlations requires a measure of the correlation between pairs of features. MdSS computed the Pearson product moment correlation between feature pairs by treating each feature as a 190-unit vector, where each unit in the feature vector corresponded to the number of participants who listed that feature for that concept in the norms (i.e., its production frequency). Features occurring in fewer than three concepts were excluded to avoid spurious correlations.

For their Experiment 3, MdSS selected a set of target features (e.g., <hunted by people>) and paired each with two concepts (e.g., *deer* and *duck*), as depicted in Figure 2. The target feature was strongly intercorrelated with the other features of one of the concepts, but weakly intercorrelated with the features of the second. Intercorrelational strength was measured by summing the shared percentage of variance between the target feature and each of the features of the concept with which it was significantly correlated. For example, according to the norms, <hunted by people> is strongly correlated with the features of *deer* (intercorrelational strength = 326), but weakly intercorrelated with the features of *duck* (intercorrelational strength = 61).

Experiment 1 extends Experiment 3a of MdSS in two ways. First, the materials are better suited to an ANOVA design because additional variables were controlled. In MdSS

Target Feature = <hunted by people>

deer

Correlated with Target Feature
<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>

Intercorrelational Strength = 326

duck

Correlated with Target Feature
<an animal>
<lives in water>
<migrates>
<swims>

Intercorrelational Strength = 61

Figure 2. Example of an item from Experiment 1. The target feature <hunted by people> was more strongly correlated with features of *deer* (strong group) than of *duck* (weak).

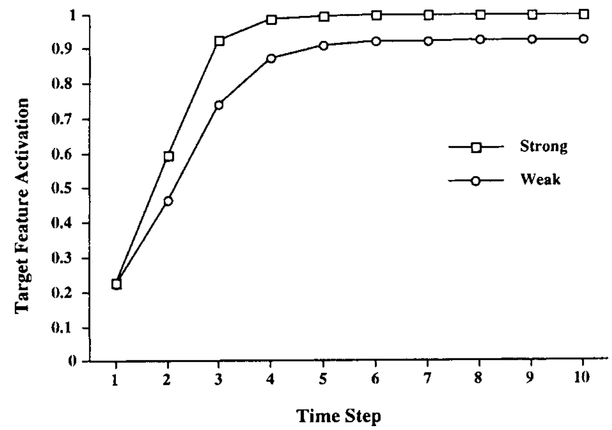

Figure 3. Simulation 1. Predictions for Experiment 1. Mean activation of the strongly and weakly intercorrelated feature units.

Experiment 3, this was not crucial because they emphasized regression analyses rather than group differences. In contrast, in both Experiments 1 and 2 herein, we focus on group differences and extensive equating of the stimuli. The second extension was to investigate the time course of the influence of correlated features by manipulating SOA (the time between the onset of the presentation of the concept and feature names).

SIMULATION 1

Predictions for Experiment 1 were taken directly from the MdSS simulation of their Experiment 3 (p. 119). Fourteen items from that experiment were used that differed in terms of the strength with which the target feature was correlated with other features of the concept, and had characteristics similar to those of our Experiment 1. Note that the feature correlation measures were similar when calculated using only the 84 concepts included in the model, as opposed to all 190 concepts from the norms. For the simulation, the word form of each concept (e.g., *deer* or *duck*) was clamped (each of its word-form units was activated to 1) and the network was allowed to settle for 10 iterations. Activation of the target feature (e.g., <hunted by people>) was recorded at each iteration. We assume that the activation of a target feature is monotonically related to the time required to verify that it is part of the concept. Five runs with independent, random starting configurations were used. Figure 3 shows that strongly intercorrelated features were more highly activated than were weakly intercorrelated features across the time course of the computation of the semantic representation of the concept. The simulation data was analysed using a two-way repeated measures ANOVA with target feature activation as the dependent variable and intercorrelational strength (strong vs. weak) and iteration (1 to 10) as the independent variables. Collapsed across the 10 iterations, feature correlations influenced target feature activation, $F(1, 13) = 10.28$.[2] Furthermore, intercorrelational strength and iteration interacted, $F(9, 117) = 5.07$. One reason for the interaction was that target feature activation for the two groups was identical after the first iteration. Because activation in the feature units was random prior to the first iteration, the network's knowledge of feature correlations that was encoded in the weights between features had not yet influenced processing. In addition, the effect of intercorrelational strength changed over further iterations; planned comparisons showed that strongly intercorrelated target features were significantly more activated for iterations 2 to 5, whereas they were marginally more activated for iterations 6 to 10 ($.06 < p < .1$). Thus, the model predicts that feature correlations will influence speeded feature verification.

The predictions for the SOA manipulation are less clear because of the uncertainty in mapping iterations in the simulation onto SOA in the human experiment, so it is probably best to take the simulation as providing a rough prediction. If earlier iterations are viewed as corresponding to a short SOA such as 300 ms and later iterations as corresponding to a long SOA such as 2,000 ms, the simulation predicts a somewhat larger effect of correlated features at the short SOA. Note that when a second ANOVA was conducted in which the feature activation for the first iteration was omitted, the interaction between intercorrelational strength and iteration remained, $F(8, 104) = 3.98$. This interaction was due primarily to ceiling effects at the later iterations. Finally, these predictions directly contrast with the view that feature correlations are not encoded in semantic memory and therefore should have no influence in Experiment 1 (Murphy & Wisniewski, 1989).

METHOD

Participants. Sixty-five University of Western Ontario

[2] In all analyses reported in this article, $p < .05$ unless otherwise noted.

TABLE 1
Manipulated and Equated Variables for Experiment 1

Factor	Strong		Weak		*t(94)*	*p*
Intercorrelational Strength	161	(54)	32	(23)	15.25	<.001
Production Frequency	12	(6)	12	(7)	-0.28	> .7
Ranked Production Frequency	8	(4)	8	(5)	-0.07	> .9
Cue Validity	0.2	(0.1)	0.2	(0.1)	0.6	> .5
Concept Familiarity	5	(1)	5	(1)	-0.44	> .6
Concept Typicality	6	(1)	6	(1)	-0.45	> .6
# Features/Concept	17	(3)	17	(3)	0.51	> .6
# Features Listed/Concept	297	(28)	297	(29)	0.04	> .9
# Letters/Concept Name	6	(2)	6	(2)	0.50	> .6
Frequency of Concept Name	19	(26)	22	(45)	-0.35	> .7

Note: Standard errors in parentheses.

undergraduate students participated for course credit, 16 per list. All participants were native speakers of English and had either normal or corrected-to-normal vision. The data for one participant was discarded because their error rate calculated over the practice, filler, and target trials was greater than 20%.

Materials. The stimuli consisted of 48 target features, each paired with 2 concepts (see Appendix A). The features of one concept were strongly intercorrelated with the target feature, whereas the features of the other were not (see Table 1 for intercorrelational strength[3].

Nine other variables that have previously been shown to influence feature verification latencies, or logically could influence them, were equated across the two groups (see Table 1). First, production frequency (the number of participants in the norming study of MdSS who listed the feature for the concept) was equated because Ashcraft (1978) and MdSS found that it predicted verification latencies. Second, ranked production frequency, the rank of the feature's production frequency in relation to the other features of the concept, was also equated because MdSS found that it predicted verification latencies. Third, cue validity of the feature (the production frequency of the feature with respect to the concept divided by the sum of the production frequencies of the feature over the concepts in which it was included) was equated across groups. Fourth, concept familiarity was equated because MdSS found that it predicted verification latencies. The familiarity measure was taken from MdSS, who asked 20 participants to judge the familiarity of the "thing that the word refers to" on a 7-point scale on which 7 corresponded to extremely familiar. Fifth, concept typicality was equated because Ashcraft found that it predicted verification latencies. The typicality measure was also taken from MdSS, who asked 20 participants to judge the typicality of the concepts with respect to the superordinates bird, mammal, fruit, vegetable, clothing, kitchen item, tool, vehicle, or weapon on a 7-point scale on which 7 corresponded to extremely typical. Sixth, the number of features listed by at least 5 of 30 participants in the MdSS norms was equated under the assumption that a concept might be activated more quickly if it contains more features. Seventh, Ashcraft found that the total number of features produced for a concept in a norming task predicted verification latency, presumably because it reflects the ease with which a concept's features can be accessed in general. Eighth, the number of letters in the concept name was equated across the strong and weak groups because word length affects reading time (Landauer & Streeter, 1973). Ninth, frequency of the concept name (Kucera & Francis, 1967) was equated because it influences reading time (Rayner & Duffy, 1986). Finally, because each feature served as a target for both groups, all variables associated with the feature itself were held constant. In summary, the strongly and weakly intercorrelated groups differed with respect to intercorrelational strength, but were equated in terms of nine potentially confounding variables, in addition to all variables associated with the features themselves.

Two lists were constructed, each containing 48 experimental concept-feature pairs, such that if a feature occurred with its strongly intercorrelated concept in List 1 (e.g., *deer* <hunted by people>), then it occurred with its weakly intercorrelated concept in List 2 (e.g., *duck* <hunted by people>), and vice versa. Thus, each list contained 24 strongly intercorrelated items and 24 weakly intercorrelated items. Each list also included 48 filler items for which the feature was not reasonably true of the concept (e.g., *buffalo* <eats seeds>). Because these items required a "no" response, they balanced the 48 experimental stimuli that required a "yes" response. The filler and experimental items were equated in two ways to avoid cueing participants to the response. First, the features were equated with respect to feature type as established by Wu and Barsalou's (1999) taxonomy. Second, the concepts for the filler trials were taken from the same superordinate categories as were the experimental concepts. The same filler items appeared in both lists. A separate set of 40 concept-feature pairs (20 requiring "yes" responses and 20 requiring "no" responses)

[3] All *t*-tests reported in this article are 2-tailed.

was constructed for the practice session. Across the practice, filler, and experimental items, no participant encountered a concept or feature name more than once.

Procedure. Participants were tested individually using PsyScope experimental software (Cohen, MacWhinney, Flatt, & Provost, 1993) on a Macintosh LC630 with a 14-inch colour Sony Trinitron monitor. They responded by pressing one of two buttons on a CMU button box that provided ms accuracy. The participants' index finger of their dominant hand was used for a "yes" response, and the index finger of their nondominant hand was used for a "no" response. Participants were randomly assigned to one of the two lists within one of the two SOA conditions (300 ms or 2,000 ms). Each trial in the 300-ms SOA condition began with a fixation point (*) in the middle of the screen for 500 ms, a blank screen for 100 ms, and then a concept name for 300 ms. After the 300-ms SOA, the feature name was presented directly below the concept name. Both the concept and feature remained on the screen until the participant responded. The ITI was 1,500 ms. The 2,000-ms SOA condition was identical except that the feature was presented 2,000 ms after the onset of the concept.

Participants were instructed to read both the concept and feature names silently, and then to indicate, as quickly and accurately as possible, whether or not the feature was reasonably true of the concept. Participants completed the 40 practice trials and then the 96 experimental trials. Trials were presented in random order. Verification latency was recorded as the time between the onset of the feature name and the participant's response. The experiment took approximately 25 minutes.

Design. ANOVAs were conducted to investigate the effects of intercorrelational strength (strong vs. weak) and SOA (300 ms vs. 2,000 ms) on verification latency and square root of the number of errors (Myers, 1979). Intercorrelational strength was within participants (F_1) and items (F_2), whereas SOA was between participants but within items. List was included as a between-participants dummy variable and item rotation group as a between-items dummy variable to stabilize variance that may result from rotating participants and items across lists (Pollatsek & Well, 1995).

RESULTS

Mean verification latency and percent errors for each condition are presented in Table 2.

Verification latencies. Trials on which an error occurred were excluded. Decision latencies greater than three standard deviations above the grand mean were replaced by the cut-off value (2%). As in Simulation 1, intercorrelational strength interacted with SOA because its influence was more pronounced for the short SOA, $F_1(1, 60) = 6.56$, $F_2(1, 46) = 3.79, p < .06$. Planned comparisons showed that with a 300-ms SOA, participants were faster to verify that a feature was part of a concept when it was strongly rather than weakly intercorrelated with the other features of the concept: 300 ms, $F_1(1, 60) = 41.65$, $F_2(1, 46) = 38.30$. This effect was less than half the magnitude, but still reliable, for the 2,000 ms, $F_1(1, 60) = 8.28$, $F_2(1, 46) = 11.89$. Collapsed across SOA, decision latencies were 59 ms faster when the feature was strongly intercorrelated with the other features of the concept (M = 853 ms, SE = 22 ms) than when it was weakly intercorrelated (M = 912 ms, SE = 23 ms), $F_1(1, 60) = 43.23$, $F_2(1, 46) = 11.73$. Finally, the 25-ms advantage for the 300-ms SOA (M = 870 ms, SE = 24 ms) versus the 2,000-ms SOA (M = 895 ms, SE = 20 ms) was significant by items, $F_2(1, 46) = 5.39$, but not by participants, $F_1 < 1$.

TABLE 2
Mean Feature Verification Latency (ms) and Percent Errors for Experiment 1

	300-ms SOA		2,000-ms SOA	
Decision Latencies				
Weak	912	(35)	913	(28)
Strong	829	(33)	876	(28)
Difference	83 *		37 *	
Percent Errors				
Weak	9.4	(1.1)	12.8	(1.7)
Strong	4.8	(1.0)	5.1	(0.7)
Difference	4.6 *		7.7 *	

Note: * indicates significant by participants and items. Standard error in parentheses.

Error rates. Intercorrelational strength and SOA did not interact, $F < 1$ in both analyses. Collapsed across SOA, participants made 4.1% fewer errors when the feature was strongly intercorrelated (M = 4.9%, SE = 0.6%) than when it was weakly intercorrelated (M = 11.1%, SE = 1.0%), $F_1(1, 60) = 34.79$, $F_2(1, 46) = 13.27$. Planned comparisons showed that participants made fewer errors for strongly intercorrelated features when the SOA was 300 ms, $F_1(1, 60) = 15.90$, $F_2(1, 46) = 15.57$, and 2,000 ms, $F_1(1, 60) = 18.98$, $F_2(1, 46) = 25.35$. Finally, error rates were nonsignificantly 1.8% lower when the SOA was 300 ms (M = 7.1%, SE = 0.8%) versus 2,000 ms (M = 8.9%, SE = 1.0%), $F_1(1, 60) = 2.77, p > .1$, $F_2(1, 46) = 1.79, p > .1$.

DISCUSSION

Experiment 1 provides additional evidence that feature correlations are encoded in semantic memory and influence performance in on-line tasks. It adds to the feature verification study of MdSS because a greater number of potentially confounding variables (including concept typicality) were equated between groups, and robust effects of feature correlations were obtained at both a short and long SOA. These results are consistent with theories of semantic memory and concept learning that incorporate a role for statistically based feature-feature relationships, such as attractor networks and the model of Billman and Heit (1988). Furthermore, assuming that the mapping from iterations in Simulation 1 to SOA in Experiment 1 is valid, the model predicted the interaction between intercorrelational strength and SOA. Note, however, that the attractor network predictions for the long SOA were not entirely accurate. The difference between the activation of the strongly and weakly intercorrelated features at later processing iterations was marginally significant, whereas the long SOA effect in the human data was robust. Nevertheless, the attractor network predictions and the Experiment 1 results directly contrast with the claim that feature correlations are not instantiated in semantic memory, and hence should not influence feature verification (Murphy & Wisniewski, 1989).

It is interesting to note that the long SOA results do not demand the notion of explicit expectancy generation. The idea that participants generate expectancies when allotted sufficient time between two stimuli is a key part of a number of theories of tasks of this sort, most notably semantic priming (Becker, 1980; Neely, 1977). However, the production frequency measures suggest that it is unlikely that participants generated the target features from either the strongly or weakly intercorrelated groups with any regularity. The average production frequency for the Experiment 1 target features was 12 of 30, indicating that they were produced for the concept by only 40% of the participants in the norming task of MdSS, even though participants produced almost 10 features per concept on average. Perhaps even more informative is the rank of the features in terms of production frequency. On average, the target features were only the eighth most likely to be generated given the concept. Finally, recall that both of these variables were equated between groups.

Experiment 2: Feature-Concept Verification

In Experiment 2, rather than pairing two concepts with one feature and presenting the concept first on each trial, two features (e.g., <is juicy> versus <is nutritious>) were paired with one concept (*grapefruit*) and the feature was presented first. In this example, the features of *grapefruit* are more strongly intercorrelated with <is juicy> (intercorrelational strength = 279) than with <is nutritious> (intercorrelational strength = 17).

SIMULATION 2

The order of presentation has direct implications for simulating the experiment and produces different predictions. Experiment 2 was simulated using the items from Simulation 1. Although these pairs were used originally to simulate a feature verification task in which the concept is presented prior to the feature (i.e., each feature is paired with two concepts), they sufficed for present purposes. Each simulated trial began by disabling the input units because the network was not trained on word-form representations corresponding to specific features. To simulate computing the feature's meaning, its corresponding unit was activated (e.g., the unit corresponding to <is juicy>). Given that each feature was represented as a single unit, activation greater than 1.0 was necessary for one of 646 feature units to significantly affect the dynamics of the system. Therefore, three independent runs were used that differed only in the degree to which the feature unit was activated (5, 10, or 15), and the results were averaged over these runs. The network was allowed to iterate five times (arbitrary). The feature unit was reactivated for each iteration because its name remained on screen during Experiment 2.

The activation levels of the other semantic features possessed by the target concept were measured at each iteration to assess how the intercorrelational strength of a single feature influences the ensuing computation of a concept containing it, and how this influence changes over time. We assumed that feature-concept verification latency is determined partly by the amount of time required for the semantic system to move from the state resulting from activating a single feature to the state representing the concept. Therefore, we assumed that verification latency in Experiment 2 is monotonically related to the number of the target concept's features that are pre-activated. Figure 4 presents the mean number of the target concept's feature units that were correctly activated for each iteration (i.e., activation > 0.5 on a 0-1 scale). A two-way repeated-measures ANOVA was conducted with the number of the target concept's feature units correctly turned on as the dependent variable and intercorrelational strength (strong vs. weak) and iteration (1 to 5) as the independent variables. Intercorrelational strength and iteration interacted in that the influence of feature correlations became more pronounced as processing progressed, $F(4, 52) = 18.72$. Furthermore, planned comparisons showed that although the effect grew over time, a significantly greater number of the target concept's feature units were correctly activated at each iteration when the initial feature was strongly, as opposed to weakly, intercorrelated with the other features of the concept. In addition, there were main effects of intercorrelational strength, $F(1, 13) = 34.36$, and iteration, $F(4, 52) = 44.66$.

Feature correlations influenced the simulation results because of the pattern completion properties of this type of network. That is, a feature activates other features in accordance with the feature-feature weights that are determined by the correlations. Therefore, a feature such as <is juicy> that is strongly intercorrelated with the other features of a concept, such as *grapefruit*, activated a greater number of the features of *grapefruit* than did a weakly intercorrelated feature such as <is nutritious>. Through

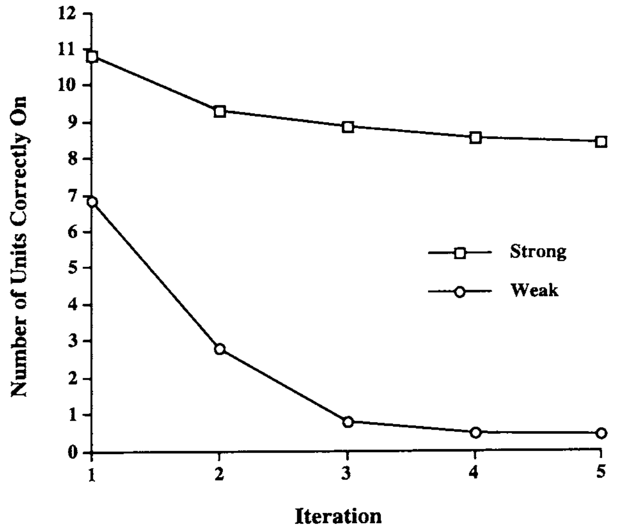

Figure 4. Simulation 2. Predictions for Experiment 2. Mean number of units correctly turned on for the strongly and weakly intercorrelated concepts.

this process of semantic pattern completion, a strongly intercorrelated feature facilitates the computation of a concept to a greater degree than does a weakly intercorrelated feature.

Predictions regarding SOA are again somewhat tenuous. Each simulated trial began by activating a semantic feature node, which corresponds to a situation in which the feature name has already been recognized. Therefore, the number of iterations in Simulation 2 is not directly comparable to Simulation 1 in which the input was a concept's word form. The general prediction that can be taken from Simulation 2, however, is that intercorrelational strength should interact with SOA in that its influence should be greater for the long SOA. The interaction occurred in Simulation 2 because the strongly intercorrelated features tend to activate a set of correlated features (a number of which are part of the target concept), and these features keep each other active over time. In contrast, the weakly intercorrelated features do not tend to be part of a cohort of this sort, so that although they initially activate other features of the target concept (because they do occur together in that concept), activation tends to spread somewhat diffusely through the network over time. Finally, the prediction of an influence of feature correlations is in direct opposition to models in which this variable plays no role.

METHOD

Participants. Sixty-four University of Western Ontario undergraduate students participated for course credit, 16 per list. All participants were native speakers of English and had either normal or corrected-to-normal vision. The data for four participants (one from each list) were discarded because their error rates calculated over the practice, filler, and target trials were greater than 20%.

Materials. Each of 36 basic-level concepts was paired with one feature that was strongly intercorrelated with its other features, and a second feature that was weakly intercorrelated (see Appendix B for the items, and Table 3 for intercorrelational strength statistics).

Eight variables were equated across groups (see Table 3). As in Experiment 1, first, production frequency and second, ranked production frequency were equated. Third, the number of concepts in which the feature appeared in the norms was equated because it may influence the degree to which a concept is activated by a feature, or predicted from it. In addition, recall that features appearing in fewer than three concepts were not included. Fourth, cue validity of the feature was again equated. Fifth, the number of words and sixth, letters per feature were equated because they influence reading time. Seventh, although we could not equate feature name frequency (Kucera & Francis, 1967), the difference favoured the weak group. Eighth, the two features paired with a concept were matched closely, but not exactly, with respect to feature type as established by Wu and Barsalou's (1999) taxonomy (see Appendix B). Finally, because each concept served as a target in both groups, all variables associated with the concept itself were held constant. In summary, the strong and weakly intercorrelated groups differed with respect to intercorrelational strength, but were equated in terms of eight potentially confounding variables, in addition to all variables associated with the concepts themselves.

Two lists were constructed, each containing 36 experimental feature-concept pairs. If a concept occurred with its strongly intercorrelated feature in List 1 (e.g., <is juicy> *grapefruit*), then it occurred with its weakly intercorrelated feature in List 2 (e.g., <is nutritious> *grapefruit*), and vice versa. Thus, each list contained 18 strongly intercorrelated items and 18 weakly intercorrelated items. Each list also included 36 filler items, for which the feature was not reasonably true of the concept, to balance the 36 experimental stimuli that required a "yes" response. As in Experiment 1, the filler and experimental items were equated in two ways to avoid cueing participants to the response: the features were equated on feature type as established by Wu and Barsalou's (1999) taxonomy; the concepts for the filler trials were taken from the same superordinate categories as were the experimental concepts. The same filler items appeared in both lists. A separate set of 40 feature-concept pairs (20 requiring "yes" responses and 20 requiring "no" responses) was constructed for the practice session. Across the practice, filler, and experimental items, no participant encountered a concept or feature name more than once.

Procedure. All aspects of the procedure were identical to Experiment 1, except that the feature name was presented prior to the concept name during each trial. Participants were instructed to read both the feature and concept names

TABLE 3
Manipulated and Equated Variables for Experiment 2

Factor	Strong		Weak		$t(70)$	p
Intercorrelational Strength	225	(76)	26	(26)	14.89	< .001
Production Frequency	11	(6)	12	(7)	−0.18	> .8
Ranked Production Frequency	8	(5)	8	(5)	−0.15	> .8
# Concepts/Feature	10	(6)	9	(6)	0.77	> .4
Cue Validity	0.2	(0.1)	0.2	(0.1)	0.17	> .8
# Words/Feature	2	(1)	2	(1)	0.21	> .8
# Letters/Feature	9	(3)	9	(3)	1.19	> .2
Summed Frequency of Feature	25	(24)	98	(134)	−3.22	< .05

Note: Standard errors in parentheses.

silently, and then to indicate, as quickly and accurately as possible, whether or not the feature was reasonably true of the concept. Participants completed the 40 practice trials and then the 72 experimental trials. Verification latency was recorded as the time between the onset of the concept name and the participant's response. The experiment took approximately 20 minutes.

Design. The design was identical to Experiment 1.

RESULTS

Mean verification latency and percent errors for each condition are presented in Table 4.

Verification latencies. As in Experiment 1, trials on which an error occurred were excluded, and decision latencies greater than three standard deviations above the grand mean were replaced by the cut-off value (1%). As in Simulation 2, intercorrelational strength interacted with SOA because its influence was more pronounced with a long SOA, $F_1(1, 56) = 3.99$; $F_2(1, 34) = 6.96$. With a 300-ms SOA, participants were faster to verify that a feature was part of a concept when it was strongly intercorrelated with the other features of the concept, $F_1(1, 56) = 6.58$, $F_2(1, 34) = 7.19$. This effect more than doubled with a 2,000-ms SOA, $F_1(1, 56) = 29.31$, $F_2(1, 34) = 41.12$. Overall, decision latencies were 56 ms faster when intercorrelational strength was strong (M = 748 ms, SE = 20 ms) than when it was weak (M = 804 ms, SE = 22 ms), $F_1(1, 56) = 31.96$, $F_2(1, 34) = 9.91$. Finally, verification latencies were 78 ms shorter when the SOA was 300 ms (M = 737 ms, SE = 17 ms) than when it was 2,000 ms (M = 815 ms, SE = 24 ms), $F_1(1, 56) = 4.15$, $F_2(1, 34) = 54.89$.

Error rates. Intercorrelational strength and SOA did not interact, $F < 1$ in both analyses. Collapsed across SOA, participants made 4.5% fewer errors when the feature was strongly intercorrelated (M = 7.4%, SE = 0.9%) than when it was weakly intercorrelated (M = 11.9%, SE = 1.0%), $F_1(1, 56) = 20.05$, $F_2(1, 34) = 5.19$. Planned comparisons showed that, with a 300-ms SOA, participants made fewer errors for strongly intercorrelated features, $F_1(1, 56) = 8.51$, $F_2(1, 34) = 10.67$. The effect was similar with a 2,000-ms SOA, $F_1(1, 56) = 11.68$, $F_2(1, 34) = 6.23$. Finally, error rates were marginally greater (2.7%) when the SOA was 300 ms (M = 11.0%, SE = 1.0%) versus 2,000-ms (M = 8.3%, SE = 1.0%), $F_1(1, 56) = 3.81, p < .06$, $F_2(1, 34) = 4.05, p < .06$.

DISCUSSION

Experiment 2 provided further evidence that feature correlations are encoded in semantic memory. Furthermore, a model incorporating this principle predicted the human results, and provided insight in terms of a feasible mechanistic explanation of the source of the effects. Specifically, the influence of intercorrelational strength arose from the pattern completion properties of this type of network. The notion of semantic pattern completion is highly plausible because this type of semantic generalization is common; people easily answer questions such as, "If something has a blade, what other features might it have?" (Sloman, 1993).

As in Experiment 1, accounting for the long SOA results does not require the notion of participants explicitly generating expectancies in the form of possible concepts given a feature name. One indication that expectancy generation played no role comes from the number of concepts in which each feature is included. On average, each strongly intercorrelated feature was part of ten concepts, and each weakly intercorrelated feature was part of nine. Therefore, it is unlikely that participants generated either the strongly or weakly intercorrelated concepts, and, if anything, they would be more likely to generate the weakly intercorrelated ones. In addition, the features did not tend to be highly associated with the concepts in that they were produced by an average of only 40% of the participants in the norming study of MdSS in which participants listed an average of 10 features per concept. In terms of ranked production frequency, the features were only the eighth most likely to be listed. Finally, the simulation mimicked the SOA by strength interaction without recourse to a mechanism by which specific concepts were explicitly generated.

General Discussion

The experiments and associated simulations provide further evidence that semantic memory includes, and perhaps

TABLE 4
Mean Feature Verification Latency (ms) and Percent Errors for Experiment 2

	300-ms SOA		2000-ms SOA	
Decision Latencies				
Weak	756	(22)	853	(36)
Strong	719	(26)	777	(31)
Difference	37 *		76 *	
Percent Errors				
Weak	13.1	(1.4)	10.7	(1.5)
Strong	8.9	(1.4)	5.9	(1.2)
Difference	4.2 *		4.8 *	

Note: * indicates significant by participants and items. Standard error in parentheses.

depends upon, statistical knowledge of feature correlations that is learned through experience with objects and entities in the everyday world. These experiments add to the list of studies demonstrating that people can learn this type of information, and provide additional evidence for theories of concept learning that include a mechanism for encoding it. The most compelling aspect of the simulations is that, assuming that the mappings from iterations to SOA are valid, the network predicted the contrasting interactions between intercorrelational strength and SOA. A greater influence of feature correlations at the short SOA was predicted and obtained when the concept name was presented first, whereas a greater influence at the long SOA was predicted and obtained when the feature name was presented first.

SPREADING ACTIVATION THEORY

Although we believe that an attractor model is the most parsimonious means of accounting for the results of Experiments 1 and 2, there are two possible ways in which spreading activation theory might account for them. The first does not rely on the notion of feature correlations. Although the extensive norming procedures of MdSS enabled equating for a large number of potentially confounding variables, there was one variable that we were not able to equate. This is what MdSS termed "feature superordinate typicality," measured as the number of concepts within a superordinate category that contain a specific feature. For example, *deer* <hunted by people> has a feature superordinate typicality score of five because five *mammal* concepts include <hunted by people>. In Experiment 1, the strongly intercorrelated concept-feature pairs had a higher average feature superordinate typicality ($M = 6.4$, $SE = 4.2$) than did the weakly intercorrelated pairs ($M = 3.4$, $SE = 2.9$), $t(94) = 4.06$. A similar pattern occurred in Experiment 2: strong ($M = 8.8$, $SE = 5.2$); weak ($M = 2.9$, $SE = 1.6$); $t(70) = 6.53$.

It is possible that this variable might influence processing in a semantic network without relying on the notion of feature correlations. Consider Collins and Loftus' (1975) model in which concept nodes are connected to other concept nodes, and to their feature nodes. When the concept name is presented first as in Experiment 1, activation might spread to the nodes representing its features, as well as to the nodes representing the other concepts that are within its superordinate category (and perhaps its superordinate category node as well). If activation spreads from each of the concept's category coordinate nodes to their feature nodes, and this activation significantly alters the state of the target feature node, then the target feature with the higher feature superordinate typicality would be more highly activated, thus accounting for the verification results. When the feature name is presented first as in Experiment 2, activation should spread from that feature node to all of the nodes that represent a concept of which it is a part. Thus, the feature with the higher superordinate typicality would activate a greater number of concept nodes from within the target concept's category, and they in turn would pass activation to the target concept. If activation spreads across the requisite links to pre-activate the target concept sufficiently, the concept from the strongly intercorrelated group would be facilitated to a greater degree.

The major problem with this account, however, is that with regards to direct concept-concept priming, research by McRae and Boisvert (1998), Moss, Ostrin, Tyler, and Marslen-Wilson (1995), and Shelton and Martin (1992) have shown that concepts must be *highly* similar to prime one another. Furthermore, Cree, McRae, and McNorgan (1999) and Lupker (1984) found that category coordinates do not prime one another unless they are highly similar, again using direct priming between concepts. Therefore, it is unlikely that indirect, mediated priming of the type described above accounts for the present results.

A second spreading activation account relies on the notion of feature correlations. Semantic network theory could be modified to include feature-feature connections, with the strength of a connection being determined by the magnitude of the correlation. This modification enables a straightforward account of our results in a spreading activation network. However, it is incumbent upon researchers from this view to propose a theory of how this knowledge is learned in this framework without resorting to the idea that these relationships must be explicitly noticed. We suspect that any spreading activation theory that incorporates these revisions would resemble an attractor model.

Finally, it is worth noting that this research does not undermine work on the knowledge-based aspects of conceptual processing (Medin, 1989; Murphy & Medin, 1985); rather, it should be regarded as complementary. A complete understanding of the various types of conceptually based tasks and the representations and computations underlying them necessitates synthesizing implicit statistically based and explicit theory-based knowledge. The work of Sloman, Love, and Ahn (1998) is an excellent example of bridging lower-level feature theories and higher-level knowledge-based theories.

Conclusions

Recent research in language development has taught us is that children are exceptionally good at extracting the structure that exists in their environment. This includes, for example, early awareness of the prosodic structure of the language (Jusczyk, Cutler, & Redanz, 1993) and phonotactic regularities (Saffran, Aslin, & Newport, 1996). Thus, language learning and use appears to hinge critically on the mind's ability to encode various sources of structure and integrate them during online comprehension and production. From this perspective, it may not be surprising that the semantic structure of objects and entities plays a key role in the computation of word meaning.

This work was supported by NSERC grant RGPIN155704 to the first author and NSERC postgraduate fellowships to the second and third authors. Part of this research formed RW's University of Western Ontario undergraduate thesis. The authors thank Michael Masson and Jeff Elman for helpful comments on earlier drafts.

Correspondence concerning this article should be addressed to Ken McRae, Department of Psychology, Social Science Centre, University of Western Ontario, London, Ontario N6A 5C2 (E-mail: mcrae@uwo.ca).

References

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

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

Billman, D. (1989). Systems of correlations in rule and category learning: Use of structured input in learning syntactic categories. *Language and Cognitive Processes, 4,* 127-155.

Billman, D., & Heit, E. (1988). Observational learning from internal feedback: A simulation of an adaptive learning method. *Cognitive Science, 12,* 587-625.

Billman, D., & Knutson, J. (1996). Unsupervised concept learning and value systematicity: A complex whole aids learning the parts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 458-475.

Clapper, J. P., & Bower, G. H. (1991). Learning and applying category knowledge in unsupervised domains. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory, Vol. 27* (pp. 65-108). Toronto, ON: Academic Press.

Cohen, J. D., MacWhinney, B., Flatt, M, & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, & Computers, 25,* 257-271.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407-423.

Cree, G. S, McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science, 23,* 371-414.

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, UK: Addison-Wesley.

Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology, 44,* 265-315.

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 features like those of two-state neurons. *Proceedings of the National Academy of Science, 81,* 3088-3092.

Jusczyk, P. W, Cutler, A., & Redanz, N. J. (1993). Infants' preference for the predominant stress patterns of English words. *Child Development, 64,* 675-687.

Kaiser, M. K., & Proffitt, D. R. (1984). The development of sensitivity to causally relevant dynamic information. *Child Development, 55,* 1614-1624.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development.* Cambridge, MA: MIT Press.

Kucera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English.* Providence, RI: Brown University Press.

Landauer, T. K., & Streeter, L. A. (1973). Structural differences between common and rare words: Failure of equivalence assumptions for theories of word recognition. *Journal of Verbal Learning and Verbal Behavior, 12,* 119-131.

Lupker, S. J. (1984). Semantic priming without association: A second look. *Journal of Verbal Learning and Verbal Behavior, 23,* 709-733.

Malt, B. C., & Smith, E. E. (1984). Correlated features in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23,* 250-269.

McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 558-572.

McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General, 126,* 99-130.

Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist, 44,* 1469-1481.

Medin, D. L., & Coley, J. D. (1998). Concepts and categorization. In J. Hochberg (Ed.), *Perception and cognition at century's end. Handbook of perception and cognition* (pp. 403-439). San Diego, CA: Academic Press.

Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 863-883.

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* (pp. 23-45). Chichester, UK: 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,* 226-254.

Pollatsek, A., & Well, A. D. (1995). On the use of counterbalanced designs in cognitive research: A suggestion for a better and more powerful analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 785-794.

Rayner, K., & Duffy, S. A. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory and Cognition, 14,* 191-201.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118,* 219-235.

Rosch, E. (1978). Principles of categorization. In E. Rosch & 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.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926-1928.

Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1191-1210.

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology, 25,* 231-280.

Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22,* 189-228.

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

Wattenmaker, W. D. (1991). Learning modes, feature correlations, and memory-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 203-222.

Wattenmaker, W. D. (1993). Incidental concept learning, feature frequency, and correlated properties. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 908-923.

Wu, L., & Barsalou, L. W. (1999). *Grounding concepts in perceptual simulation: I. Evidence from property generation.* Unpublished manuscript.

Younger, B. (1990). Infants' detection of correlations among feature categories. *Child Development, 6,* 614-620.

Younger, B. A., & Cohen, L. B. (1983). Infant perception of correlations among attributes. *Child Development, 54,* 858-867.

Younger, B. A., & Cohen, L. B. (1986). Developmental change in infants' perception of correlations among attributes. *Child Development, 57,* 803-815.

Appendix A

Stimuli for Experiment 1

Feature	Concept — Strong	Concept — Weak
found in kitchens	toaster	sink
has a long tail	rat	pony
has a seat	tricycle	chair
has eyes	fawn	hawk
has leaves	lettuce	pineapple
has legs	pony	stork
has teeth	lion	rat
has wheels	bus	cannon
hunted by people	deer	duck
is breakable	bottle	crayon
is dangerous	rifle	motorcycle
is electrical	microwave	drill
is grey	rock	mouse
is hand held	pistol	screwdriver
is nutritious	cauliflower	grapefruit
is round	cherry	peas
is sharp	spear	fork

lives in herds	caribou	sheep
migrates	duck	caribou
used for juice	grapefruit	grape
worn by women	dress	nylons
made of leather	boots	belt
is colourful	budgie	carpet
used for storage	closet	shelves
has an engine	dunebuggy	yacht
has a handle	axe	cup
has doors	cupboard	car
has drawers	dresser	desk
has patterns	plate	mug
has seeds	orange	cucumber
is brown	moose	coconut
used for transportation	van	horse
is comfortable	cushion	slippers
is crunchy	radish	apple
is edible	cauliflower	chicken
is green	cucumber	lime
is juicy	lime	beets
is orange	tangerine	carrot
is pointed	spear	screws
is rectangular	freezer	dresser
is transparent	jar	nylons
is tropical	coconut	parakeet
worn for warmth	boots	shirt
made of plastic	cup	fork
made of steel	knife	pliers
used for protection	pistol	dog
used in circuses	lion	cannon
used for war	gun	sword

Appendix B

Stimuli for Experiment 2, Along with Feature Type According to Taxonomy of Wu and Barsalou (1999)

Concept	Feature — Strong	Type	Feature — Weak	Type
beans	is nutritious	e sys	is yellow	e se
blouse	made of cotton	made of	is soft	e se
bottle	has a lid	e ce	has a neck	e ce
buzzard	has wings	e ce	is black	e se
cabbage	grows in gardens	location	is round	e se
canary	chirps	e b	has 2 legs	e q
car	has 4 wheels	e q	is expensive	e sys
caribou	is herbivorous	superordinate	migrates	e b
carrot	is crunchy	e si	is orange	e se
chair	has armrests	e ce	has a seat	e ce
cherry	is sweet	e si	is red	e se
couch	used for relaxing	function	made of leather	made of
crow	flies	e b	is loud	e b
cup	made of china	made of	made of glass	made of
dagger	is sharp	e se	is pointed	e se

grapefruit	is juicy	e si	is nutritious	e sys
gun	is dangerous	e sys	is black	e se
missile	used for killing	function	used by the army	participant
moose	has hooves	e ce	has hair	e ce
motorcycle	has an engine	e ci	is dangerous	e sys
mouse	has a tail	e ce	is grey	e se
ostrich	has feathers	e ce	has legs	e ce
parakeet	has a beak	e ce	is tropical	e sys
pig	has 4 legs	e q	is dirty	e se
pineapple	is juicy	e si	is yellow	e se
pumpkin	has seeds	e ci	is orange	e se
rat	has whiskers	e ce	has teeth	e ci
sandals	has soles	e ce	has straps	e ce
slingshot	used for shooting	function	used by children	participant
spear	is sharp	e se	is thin	e se
stone	is smooth	e se	is cold	e se
stork	flies	e b	is white	e se
tricycle	has wheels	e ce	used by children	participant
trousers	has pockets	e ce	is comfortable	evaluation
van	has wheels	e ce	has doors	e ce
vulture	has wings	e ce	has claws	e ce

Legend: e ce = entity external component; e ci = entity internal component; e se = entity external surface property; e si = entity internal surface property; e b = entity behaviour; e q = entity quantity; e sys = entity systemic property (includes "abstract properties" of an object as a whole)

Sommaire

Le rôle des corrélations de caractéristiques dans la mémoire sémantique est un problème central dans la représentation conceptuelle. Dans deux versions de la vérification des caractéristiques, les sujets ont vérifié plus rapidement qu'une corrélation (<est juteux>) fait partie d'un concept (*pamplemousse*) si elle est plutôt fortement que faiblement corrélée à d'autres caractéristiques de ce concept. Des interactions contrastantes entre les corrélations de caractéristiques et des apparitions-stimulus asynchrones (SOA) ont été découvertes quand le concept était présenté avant la caractéristique. Un modèle de réseau attracteur sur la signification des mots, qui apprend et utilise naturellement les corrélations entre les caractéristiques, a prédit ces interasctions. La présente recherche prouve davantage que la mémoire sémantique comprend une connaissance statistique implicitement acquise des futures relations, en contraste avec des théories comme la propagation des réseaux d'activation, où les corrélations des caractéristiques ne jouent aucun rôle.