

Shared Features Dominate The Number-Of-Features Effect

Ray Grondin (rgrondin@uwo.ca)
Stephen J. Lupker (lupker@uwo.ca)
Ken McRae (mcrae@uwo.ca)

Department of Psychology, Social Science Centre
University of Western Ontario
London, Ontario, Canada, N6A5C2

Abstract

When asked to list semantic features for concrete concepts, participants list many features for some concepts and few for others. Concepts with many semantic features have been reported to be processed faster in lexical decision, naming, and semantic decision tasks (Pexman, Holyk, & Monfils, 2003; Pexman, Lupker, & Hino, 2002). Using a much larger and better controlled set of items, we replicated the number-of-features (NoF) effect in both lexical and semantic decision (Experiment 1). We then investigated the relationship between NoF and feature distinctiveness. Shared features are those which appear in many concepts (<has four legs>) whereas distinctive features appear in few concepts (<moos>). Keeping total NoF constant, decision latencies were shorter for concepts with many shared features versus those with few shared features in lexical and semantic decision, with a larger difference obtaining in semantic decision (Experiment 2). Manipulating shared or distinctive features to create low versus high levels of NoF revealed a larger NoF advantage for concepts with many shared features than for those with many distinctive features (Experiment 3). It is concluded that shared features play a dominant role in the NoF effect, at least in lexical and semantic decision tasks.

Keywords: semantic categorization; lexical decision; word recognition; semantic features; shared features; distinctive features

Introduction

People use language to convey messages, and inherent in our ability to understand these messages is our ability to compute the meaning of individual words. The goal of the current research is to further our understanding of the computation of word meaning. In particular, we investigate an emerging finding that the 'richness' of a word's semantic representation influences performance in speeded tasks involving the computation of its meaning.

One example of such a result is the ambiguity advantage (Hino & Lupker, 1996). Specifically, words with multiple meanings (*bowl*) are responded to faster than words with a single meaning (*tent*) in tasks such as lexical decision (*Is the letter string an English word?*) and naming (*Read the presented word aloud*). Words with multiple meanings are assumed to have richer semantic representations because multiple instead of single meanings have to be encoded.

Similarly, words that refer to concrete objects (*robin*) are responded to faster than words that refer to abstract concepts (*justice*). Again, this is true in both lexical decision (Binder, Westbury, McKiernan, Possing, & Medler, 2005) and

naming (Strain, Patterson, & Seidenberg, 1995). A number of researchers argue that this difference can be explained in terms of a richer semantic representation for concrete words. For example, Paivio (1986) claimed that in addition to being able to verbally reason about both concrete and abstract things, people can also generate mental iconic images for concrete words because they refer to physical things in the world which we can perceive. He argued that this additional information associated with concrete words makes their mental representations richer and easier to process.

Plaut and Shallice (1993) approached this issue using a feature-based representation of word meaning. They hypothesized that a major difference between concrete and abstract words is the number of features. That is, although we can easily generate many features for concrete entities and objects (a *robin* <has wings>, <flies>, <eats worms>, <has a red breast>, etc.), it is much harder to generate features for abstract words. Plaut and Shallice reported that patients with deep dyslexia make more errors when reading abstract words than when reading concrete words. Using a distributed representation of word meaning where concrete words had on average more features than abstract words, they simulated deep dyslexia in a connectionist network by randomly removing connections between and within layers of the network. They found that because concrete concepts had more features in the model and thus generated stronger attractors than did abstract concepts, concrete words were often less susceptible to network damage.

All of these explanations rest on the assumption that the underlying representations differ depending on number of meanings and concreteness. Although this assumption may be correct, to test the richness hypothesis directly, it would be better to have representations of word meaning that are generated (as directly as possible) from people's actual conceptual representations.

McRae, Cree, Seidenberg, and McNorgan (2005) presented participants with living (*robin*) and nonliving (*chair*) thing concepts and had them list descriptive features for each. For example, for *robin*, participants listed features such as those presented above. Of course, people cannot introspectively tell us everything that exists in their conceptual representations, but it is assumed that what they do tell us provides a reasonable window into those representations (Medin, 1989).

McRae et al. (2005) used 725 participants to collect these semantic feature production norms for 541 concepts. This large set was used to define a semantic space

consisting of 2526 featural dimensions, and enabled the calculation of many statistics such as correlations between features and feature distinctiveness. In conjunction with various other measures (word frequency, word length, conceptual familiarity, orthographic and phonological neighborhoods), these empirically-derived conceptual representations provide a rich basis for testing theories of semantic representation and computation.

For our purposes, one advantage of the concepts contained in McRae et al.'s (2005) norms is that each refers to a concrete object (living or nonliving) and each has, as much as possible, a single meaning. Although facilitation has been obtained for multiple over single meaning words and for concrete over abstract words, it is not clear whether this facilitation is due to the difference in richness or to some other confounding variable. For instance, it is possible that words with multiple meanings are processed more quickly not only because their representations are richer, but also because people have thought about them more deeply during learning because it is necessary to tease apart their multiple meanings. Also, because of the way that we interact with concrete objects but not with abstract concepts, it is likely that concrete representations span different parts of the brain (sensory and motor). On the other hand, if it really is a difference in semantic richness that is underlying these facilitation effects, then we expect to find similar results when comparing words that differ in semantic richness within the same word type, specifically, the single meaning concrete nouns found in McRae et al.'s norms.

Pexman and colleagues selected stimuli from McRae, de Sa, and Seidenberg's (1997) norms, which included 190 of the 541 concepts found in McRae et al.'s (2005) norms. They began their investigation by generating two sets of concepts. One set contained 25 low number-of-features (low NoF) concepts, and the other contained 25 high number-of-features (high NoF) concepts. For their lexical decision tasks, they also generated two sets of 50 nonword filler items. The first contained pronounceable pseudowords whose spelling and sound do not correspond to any English word (*merod*). The second contained pseudohomophones whose spelling does not correspond to any English word but whose sound does (*keap*). In their first study, they combined the low and high NoF concepts with the pseudoword fillers and found that lexical decision latencies were shorter for high than for low NoF concepts. This effect was even larger when the fillers were pseudohomophones. They reported similar results in a subsequent study in which they asked participants to name the same low and high NoF items aloud (Pexman, Lupker, & Hino, 2002).

Pexman, Holyk, and MonFils (2003) used all 190 concepts as well as 190 filler abstract concepts. They had participants perform a concreteness decision task in which participants decided whether each word referred to a concrete object or to something abstract. Whereas previous studies had found processing differences *between* concrete and abstract concepts, Pexman et al. found semantic richness (NoF) effects *among* the concrete concepts.

Experiment 1

A closer investigation of the items used in Pexman et al.'s (2002; 2003) studies, however, reveals that some variables known to influence word processing (word frequency and word length) were, to some degree, confounded with the number-of-features manipulation. Although Pexman and colleagues addressed this issue by partialling out the influence of these variables using multiple regression, it is possible that the observed NoF effects were due to the combined influence of these confounded variables. Therefore, the purpose of Experiment 1 is to test whether Pexman et al.'s results replicate in both lexical and concreteness decision tasks. It was possible to construct larger lists of concepts that are better balanced on more variables because we had access to a larger set of norms (541 instead of 190 concepts).

Method

Participants. Thirty-four Psychology undergraduate students at the University of Western Ontario participated for course credit. Seventeen were assigned to lexical decision and 17 to concreteness decision. In all Experiments reported herein, all participants had normal or corrected-to-normal visual acuity, and were native English speakers.

Materials. Two sets of target words referring to concrete objects were generated from McRae et al.'s (2005) semantic feature production norms. One set consisted of 64 low NoF concepts and the other consisted of 64 high NoF concepts. The two sets were matched carefully on a number of potentially confounding variables (Table 1). These included word frequency, which was computed using the natural logarithm of the singular plus plural counts taken from the British National Corpus (BNC) online search engine (Burnard, 2000). Concept familiarity was measured by asking 20 participants to rate, on a 9-point scale, with 1 corresponding to *not at all familiar*, and 9 corresponding to *highly familiar*, *'How familiar are you with the thing the word refers to?'* Number of letters, number of phonemes, number of syllables, and orthographic neighborhood size (Coltheart, Davelaar, Jonasson, & Besner, 1977) were all computed using the N-watch program (Davis, 2005). Semantic density, a measure of the degree to which a concept's features are intercorrelated, was calculated from McRae et al.'s norms. In the norms, each feature is a vector of production frequencies (number of participants listing that feature) across the 541 concepts. Proportion of shared variance for each feature-vector pair was calculated by squaring the correlation between the two vectors. Semantic density for a concept is the sum of the proportion of shared variances for each pair of features that are included in that concept. Finally, because the extent to which different types of concrete objects are processed differentially is unclear (Laws & Gale, 2002), we also matched the groups according to the following category breakdown: creatures, fruits and vegetables, and nonliving things.

Table 1: Characteristics of Experiment 1 Stimuli

Variable	Low NoF		High NoF	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Number of features (NoF)	9.0	0.2	15.7	0.3
ln(BNC) frequency	6.4	0.2	6.4	0.2
Familiarity	5.7	0.2	5.7	0.2
Number of letters	5.4	0.2	5.4	0.2
Number of phonemes	4.4	0.2	4.5	0.2
Number of syllables	1.6	0.1	1.7	0.1
Orth. neighborhood size (N)	4.6	0.8	4.4	0.7
Semantic density	156.9	20.1	155.3	11.3
Number of creatures	18	—	21	—
Number of fruits/vegetables	11	—	8	—
Number of nonliving things	35	—	35	—

Note. NoF = Number of Features, ln = the natural logarithm (log_e), BNC = British National Corpus

Lexical decision filler items consisted of 128 pronounceable pseudowords and concreteness ("semantic") decision filler items consisted of 128 abstract concepts. Both sets were matched with the target items on the mean number of letters.

Procedure. Participants were tested individually using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993) on a Macintosh computer equipped with a CMU button box. Letters were approximately 0.5 cm high, black, and presented on a white background. One item was presented at a time and participants made either a lexical or semantic decision depending on which task they were assigned. Participants used the index finger of their dominant hand for a 'yes' response and the index finger of their nondominant hand for a 'no' response. Decision latencies were measured from the onset of the stimulus presentation to the onset of the button press. Items were presented until the participant made a decision and were presented in a different random order for each participant. Participants were instructed to make their decisions as quickly and accurately as possible.

Results and Discussion

Separate subject (t_1) and item (t_2) analyses were performed on decision latencies for concrete concepts.¹ Errors (lexical decision: 4.6% of trials; semantic decision: 3.9%) were removed from the analyses and correct decisions that exceeded 3 standard deviations above the grand mean for the target words were replaced with the cutoff value (lexical decision: 1.8%; semantic decision: 1.4%). The independent variable was NoF (low versus high) which was within-subjects and between-items.

¹ Error analyses were performed for all Experiments, but the differences between conditions were generally small and where significant differences were observed, there was no speed-accuracy tradeoff. Therefore, error analyses are not presented.

Lexical decision latencies to high NoF concepts ($M = 593$ ms, $SE = 20$ ms) were 30 ms shorter than to low NoF concepts ($M = 623$ ms, $SE = 19$ ms), $t_1(16) = 7.37$, $p < .001$, $t_2(126) = 2.34$, $p < .05$. Semantic decision latencies to high NoF concepts ($M = 637$ ms, $SE = 14$ ms) were 29 ms shorter than to low NoF concepts ($M = 666$ ms, $SE = 16$ ms), $t_1(16) = 4.27$, $p < .01$, $t_2(126) = 2.51$, $p < .05$. Thus, using this tightly controlled and larger set of items, we replicated Pexman et al.'s (2002; 2003) number-of-features effect. Words rich in semantic representation (as measured by the number of features listed in the norms) were responded to faster than words that are less rich.

In the remainder of this article, we investigate a potential source of the NoF effect by contrasting *shared* versus *distinctive* features. Shared features are those that occur in many concepts (<has four legs> and <is hard>) whereas distinctive features are those that occur in few concepts (<moos> and <oinks>). Shared features denote commonalities among concepts, and thus indicate ways in which concepts are similar to one another. In contrast, distinctive features denote differences, and thus help people to discriminate among concepts.

The relative contribution of these two feature types to processing appears to be task dependant. For instance, Humphreys, Riddoch, and Quinlan (1988) found that people were faster to name pictures of objects belonging to categories whose exemplars were structurally dissimilar (clothing and furniture) than pictures of objects from structurally similar categories (insects, fruits, and vegetables). However, when asking participants to make category decisions involving these same stimuli, Riddoch and Humphreys (1987) found that they could do so more quickly when the pictures of objects belonged to categories whose exemplars were structurally similar than when they were dissimilar. Thus, distinctive features appear to facilitate processing when the task requires distinguishing an item from among similar items (picture naming), whereas shared features appear to facilitate processing when the item has to be identified as a member of a larger category.

In the cases of deciding whether a string of letters is a word, or whether it refers to a concrete object, it is possible, at least theoretically, that people can initiate their response prior to precisely identifying or distinguishing a concept from among similar concepts. As such, it seems reasonable to predict that shared features contribute more than distinctive features to the processing advantage found in the number-of-features effect. That is, the more shared features a concept has, the more likely it is to be processed quickly in lexical and semantic decision tasks.

Cree and McRae (2003) defined a feature as shared if it was listed for more than 2 of the 541 concepts and distinctive if it was listed for only 1 or 2 concepts. They also computed feature distinctiveness as a continuous dimension (the multiplicative inverse of the number of concepts in which a feature occurred), but for present purposes, we focus on the shared versus distinctive binary measure.

In Experiment 1 and Pexman et al.'s (2002) studies, both shared and distinctive features were higher for the high NoF concepts. Therefore, these Experiments provide no insight into the relative contributions of shared versus distinctive features. One way to test the relative contributions is to directly contrast the number of shared versus distinctive features while holding NoF constant (Experiment 2). Another way is to create the NoF manipulation by altering either the number of shared or distinctive features while holding the other constant (Experiment 3).

Experiment 2

The purpose of Experiment 2 is to investigate whether lexical and semantic decisions are systematically influenced when the number of shared (and distinctive) features is manipulated while holding NoF constant.

Method

Participants. Forty-nine undergraduate students at the University of Western Ontario received \$10 for their participation. Twenty-five were assigned to lexical decision and 24 to semantic decision.

Materials. Two sets of target words referring to concrete objects were generated from McRae et al.'s (2005) norms. One set consisted of 55 low number-of-shared-features concepts and the other consisted of 55 high number-of-shared-features concepts. The two sets were tightly matched on the same variables described in Experiment 1, plus NoF.

Lexical decision filler items consisted of 110 pronounceable pseudowords and semantic decision filler items consisted of 110 abstract concepts. Both sets were matched with the target items on the mean number of letters.

Procedure. The procedure was identical to Experiment 1.

Results and Discussion

Errors (lexical decision: 3.6%; semantic decision: 3.8%) again were removed from the analyses and correct decisions that exceeded 3 standard deviations above the grand mean for the target concepts were replaced with the cutoff value (1.6% of trials for both lexical and semantic decision). The independent variable was number-of-shared-features (low versus high) which was within-subjects and between-items.

Lexical decision latencies to concepts with a high number-of-shared-features ($M = 544$ ms, $SE = 12$ ms) were 11 ms shorter than to those with a low number-of-shared-features ($M = 555$ ms, $SE = 13$ ms), which was significant by subjects, $t_1(24) = 4.60$, $p < .001$, but not by items, $t_2(108) = 1.07$, $p > .2$. Semantic decision latencies to concepts with a high number-of-shared-features ($M = 714$ ms, $SE = 29$ ms) were 42 ms shorter than to those with a low number-of-shared-features ($M = 756$ ms, $SE = 31$ ms), $t_1(23) = 6.17$, $p < .001$, $t_2(108) = 2.18$, $p < .05$.

Thus, increasing the number of shared features facilitates both lexical and semantic decision, although the degree of facilitation is greater for semantic decision. Obtaining a

larger effect in semantic decision is not particularly surprising. Although there was no difference in effect size between lexical (30 ms) and semantic (29 ms) decision in Experiment 1, a number of studies have found stronger effects of semantic manipulations on semantic than lexical decision tasks (McRae & Boisvert, 1998; Becker, Moscovitch, Behrmann, & Joordens, 1997). Although it is clear that participants must compute the meaning of a word to decide whether it refers to something that is concrete or abstract, computation of meaning may be less strongly related to making a lexical decision (Pexman et al., 2002). That is, it appears that lexical decisions can be made on the basis of some combination of orthographic, phonological, and semantic knowledge.

If adding shared features facilitates processing, what about adding distinctive features? Under the richness hypothesis, one might predict that adding features, regardless of type, facilitates processing to some degree because the result is a richer semantic representation. However, it is also possible that adding distinctive features inhibits processing because as distinctive features are added to a concept, the concept becomes more dissimilar to other concepts. This could be particularly true in a task such as semantic verification in which latencies are longest for concepts that are dissimilar to other category members (atypical; Rips, Shoben, & Smith, 1973). In fact, although shared features were manipulated in Experiment 2, if we reverse the logic and focus on distinctive rather than shared features, there is evidence that distinctive features inhibit processing in both tasks. That is, since each feature is either shared or distinctive, the numbers of shared and distinctive features are simply the complements of one another in Experiment 2. Re-labeling the low-shared condition as high-distinctive and the high-shared condition as low-distinctive suggests that decision latencies are longer to concepts with many versus few distinctive features. Although this interpretation is possible, it is unlikely given that in Experiment 1 we added, on average, more distinctive features (3.9) than shared features (2.9) in the high NoF condition and facilitation was obtained.

Experiment 3

Experiment 3 is designed to overcome the ambiguity in interpreting Experiment 2 by systematically testing the influence of shared and distinctive features while holding the other constant in both lexical and semantic decision.

Method

Participants. Eighty-nine undergraduate Psychology students at the University of Western Ontario received course credit for their participation. Forty-five were assigned to lexical decision and 44 to semantic decision.

Materials. Four sets of 20 words referring to concrete objects were generated. In the first two sets, the number-of-distinctive-features was held constant while the number-of-shared-features was manipulated. In the other two sets, the number-of-shared-features was held constant and the number-of-distinctive-features was manipulated. Again,

these four sets were tightly matched on the same variables described in Experiment 1.

Lexical decision filler items consisted of 80 pronounceable pseudowords and semantic decision filler items consisted of 80 abstract concepts. Both sets were matched with the target items on the mean number of letters.

Procedure. The procedure was identical to Experiment 1.

Results and Discussion

Errors (lexical decision: 3.3%; semantic decision: 4.7%) were removed from the analyses and correct decisions that exceeded 3 standard deviations above the grand mean for the target concepts were replaced with the cutoff value (1.7% of trials for both lexical and semantic decision). The independent variables were type of manipulated feature (shared versus distinctive) and NoF (low versus high), both of which were within-subjects and between-items. Mean decision latencies are presented in Table 2.

In lexical decision, feature type interacted with NoF by subjects, $F_1(1, 44) = 7.59, p < .01$, but not by items, $F_2(1, 76) = 1.27, p > .2$. Planned comparisons revealed that decision latencies to high NoF concepts were marginally shorter than to low NoF concepts when shared features were manipulated, $F_1(1, 86) = 13.19, p < .05, F_2(1, 76) = 1.76, p > .1$. There was no NoF difference when distinctive features were manipulated, $F_1(1, 86) < 1, F_2(1, 76) < 1$.

Decision latencies were 10 ms shorter for high NoF ($M = 598$ ms, $SE = 9$ ms) than for low NoF concepts ($M = 608$ ms, $SE = 9$ ms) overall, which was significant by subjects, $F_1(1, 44) = 5.65, p < .05$, but not by items, $F_2(1, 76) < 1$. Decision latencies were 10 ms shorter for concepts with distinctive features manipulated ($M = 598$ ms, $SE = 8$ ms) than for those with shared features manipulated ($M = 608$ ms, $SE = 9$ ms), which was significant by subjects, $F_1(1, 44) = 6.37, p < .05$, but not by items, $F_2(1, 76) < 1$.

In semantic decision, the interaction between manipulated feature type and NoF was significant by subjects, $F_1(1, 43) = 15.98, p < .001$, and marginal by items, $F_2(1, 76) = 2.54, p > .1$. Planned comparisons revealed that, when number of shared features was manipulated, decision latencies to high NoF concepts were shorter than to low NoF concepts, $F_1(1, 85) = 63.03, p < .01, F_2(1, 76) = 9.59, p < .01$. However, when number of distinctive features was manipulated, the NoF effect was significant by subjects, $F_1(1, 85) = 4.65, p < .05$, but not by items, $F_2(1, 76) < 1$.

Decision latencies were 51 ms shorter for high NoF concepts ($M = 713$ ms, $SE = 13$ ms) than for low NoF concepts ($M = 764$ ms, $SE = 14$ ms), $F_1(1, 43) = 53.59, p < .001, F_2(1, 76) = 7.67, p < .01$. Decision latencies were 27 ms shorter for concepts with distinctive features manipulated ($M = 725$ ms, $SE = 13$ ms) than for concepts with shared features manipulated ($M = 752$ ms, $SE = 14$ ms), which was significant by subjects, $F_1(1, 43) = 26.16, p < .001$, but not by items, $F_2(1, 76) = 2.12, p > .1$.

Table 2: Decision latencies (ms) for Experiment 3

	Manipulated Feature Type			
	Shared		Distinctive	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Lexical Decision				
Low NoF	621	13	596	12
High NoF	596	14	600	12
Difference	25 *		-4	
Semantic Decision				
Low NoF	792	19	736	19
High NoF	711	18	714	19
Difference	81 **		22 *	

Note. NoF = Number of Features, * = significant by subjects, ** = significant by subjects and items

As in Experiment 2, in Experiment 3 increasing the number of shared features marginally decreased lexical decision latencies and significantly decreased semantic decision latencies. It was also the case that increasing the number of distinctive features did not decrease lexical decision latencies, but did marginally decrease semantic decision latencies although the effect was smaller than for shared features. Thus, while adding distinctive features does not inhibit performance, Experiment 3 confirmed that adding shared features provides more facilitation than adding distinctive features in lexical and semantic decision.

General Discussion

The present experiments demonstrate that the richness of a word's semantic representation, in terms of the number of features, influences speeded decisions involving the computation of its meaning. Using empirically-derived feature lists, Pexman and colleagues found that concepts with many features were responded to faster than concepts with few features, and this was taken as evidence of a processing advantage for concepts rich in semantic representation. We extended this research in two ways. First, Experiment 1 showed that the effect is robust when using large lists of items that are tightly controlled on many variables. Second, Experiments 2 and 3 showed that increasing the number of shared features facilitates processing to a greater extent than does increasing the number of distinctive features.

Why do shared features provide more facilitation, particularly in the semantic (concreteness) decision task? We computed the degree to which features are shared with respect to 541 concrete concepts. Thus, in terms of the present concreteness decision task, the concepts that possess numerous shared features are precisely the ones that are most similar to other concrete objects. That is, concepts with many shared features are more typical members

(assuming equal number of distinctive features) of the category 'concrete objects' and so are categorized fastest.

The fact that adding distinctive features does not inhibit concreteness decisions is interesting because it seems to conflict with a traditional typicality account. However, in a broad and variable category such as 'concrete objects', it is not clear that adding distinctive features necessarily makes a concept less typical of concrete objects. For instance, a feature like <moos> is distinctive of the concept *cow*, but retrieving this feature would presumably not inhibit categorizing *cow* as a concrete object.

The relative contribution of these two feature types is likely to be task specific. Recall that Riddoch and Humphreys (1987) and Humphreys et al. (1988) found that whereas pictures of exemplars taken from categories composed of numerous structurally similar items (shared features) were categorized faster than those from structurally distinct categories, the opposite pattern was found when the pictures were named. This suggests that distinctive features matter more than shared features in tasks such as picture naming where distinctive features are crucial to distinguishing an item from among similar items. In fact, preliminary results support this prediction (P. Pexman, personal communication, Nov, 2005).

One implication of the number-of-features effects and the differential role of shared and distinctive features in processing is that they provide strong evidence against a static, localist view of conceptual representation. Instead, concepts appear to be highly distributed and dynamically generated depending on the task. Thus, these results have important implications for understanding and constructing models of the computation of word meaning.

Acknowledgments

This work was supported by a NSERC Master's Scholarship and a SHARCNET Graduate Research Fellowship to RG, and NSERC grant OGP0155704 and NIH grant MH6051701 to KM. We also thank Samantha Proffitt for her assistance with the experiments.

References

Becker, S., Moscovitch, M., Behrmann, M., & Joordens, S. (1997). Long-term semantic priming: A computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(5), 1059-1082.

Binder, J. R., Westbury, C. F., McKiernan, K. A., Possing, E. T., & Medler, D. A. (2005). Distinct brain systems for processing concrete and abstract concepts. *Journal of Cognitive Neuroscience*, 17, 1-13.

Burnard, L. (2000). *British National Corpus user reference guide Version 2.0*. Retrieved May 20, 2002, from British National Corpus, Oxford University Computing Service, <http://hcu.ox.ac.uk/BNC/World/HTML/urg.html>.

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

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and performance* (Vol. 6, pp. 535-555). Hillsdale, NJ: Erlbaum.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of *chipmunk*, *cherry*, *chisel*, *cheese*, and *cello* (and many other such concrete nouns). *Journal of Experimental Psychology: General*, 132, 163-201.

Davis, C. J. (2005). N-Watch: A program for deriving neighborhood size and other psycholinguistic statistics. *Behavior Research Methods*, 37, 65-70.

Hino, Y., & Lupker, S. J. (1996). Effects of polysemy in lexical decision and naming: An alternative to lexical access accounts. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 1331-1356.

Humphreys, G. W., Riddoch, M. J., & Quinlan, P. T. (1988). Cascade processes in picture identification. *Cognitive Neuropsychology*, 5, 67-103.

Laws, K. R., & Gale, T. M. (2002). Category-specific naming and the visual characteristics of line drawn stimuli. *Cortex*, 38, 7-21.

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

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37, 547-559.

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, 1269-1481.

Paivio, A. (1986). *Mental representations: A Dual coding approach*. New York: Oxford University Press.

Pexman, P. M., Holyk, G. G., & MonFils, M-H. (2003). Number-of-features effects and semantic processing. *Memory & Cognition*, 31, 842-855.

Pexman, P. M., Lupker, S. J., & Hino, Y. (2002). The impact of feedback semantics in visual word recognition: Number-of-features effects in lexical decision and naming tasks. *Psychonomic Bulletin & Review*, 9, 542-549.

Plaut, D.C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10, 377-500.

Riddoch, M. J., & Humphreys, G. W. (1987). Picture naming. In G. W. Humphreys & M. J. Riddoch (Eds.), *Visual object processing: A cognitive neuropsychological approach* (pp. 107-143). London: Erlbaum.

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.

Strain, E., Patterson, K., Seidenberg, M. S. (1995) Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1140-1154.