Using a High-dimensional Memory Model to Evaluate the Properties of Abstract and Concrete Words

Chad Audet (chad@cassandra.ucr.edu)
Department of Psychology; University of California, Riverside
and Language Machines, Inc.
Riverside, CA 92521

Curt Burgess (curt@doumi.ucr.edu)
Department of Psychology; University of California, Riverside
and Language Machines, Inc.
Riverside, CA 92521

Abstract

The evidence that the comprehension of abstract and concrete words differ prompts one to consider how the lexical representations for these word types differ. The context-availability model (Schwanenflugel & Shoben, 1983) suggests that abstract words are more difficult to process because associated contextual information stored in memory for these words is more difficult to retrieve than for concrete words. Schwanenflugel (1991) provides two hypotheses regarding how these differences in retrieval of contextual information may come about. Three simulations using context representations from the Hyperspace Analogue to Language (HAL) model of memory (Burgess & Lund, 1997; Lund & Burgess, 1996) are used to evaluate Schwanenflugel's hypotheses, as well as to provide insight into the representational differences between abstract and concrete words.

While most empirical results have consistently demonstrated that abstract words like "permission" and "issue" are more difficult to comprehend than concrete words like "newspaper" and "apple," the source of this relative difficulty for understanding abstract words is unclear. The finding that concrete words are more readily processed by language users--the "concreteness effect"--has been found in a number of experiments using a variety of tasks (for reviews, see Balota, Ferraro, & Conner, 1991; and Schwanenflugel, 1991; for representative studies that have failed to find concreteness effects, see these same reviews). Concreteness effects have been shown for words presented in lexical decision (Bleasdale, 1987; Chiarello, Senehi, & Nuding, 1987; Ransdell & Fischler, 1987; Schwanenflugel, Harnishfeger, & Stowe, 1988; Schwanenflugel & Shoben, 1983), naming (Bleasdale, 1987; de Groot, 1989; Schwanenflugel & Stowe, 1989), free recall (Paivio, 1986; Ransdell & Fischler, 1987; Schwanenflugel, Akin, & Luh, 1992), and word association (de Groot, 1989). Concreteness effects have also been seen in experimental tasks that involved comprehension or recall of sentences or paragraphs that differed on concreteness (Belmore, Yates, Bellack, Jones, & Rosenquist, 1982; Holmes & Langford, 1976; Ransdell & Fischler, 1989; Schwanenflugel & Shoben, 1983).

Given the wealth of evidence that suggests that the processing of concrete and abstract words differs, an explanation of the representational differences between these word types would seem to be essential for any robust theory of word meaning. One potential explanation of this difference is extended by the context-availability model discussed by Schwanenflugel and Shoben (1983; and Schwanenflugel, 1991). This model posits that language comprehension is aided by the retrieval from memory of contextual information associated with the material being processed. If the appropriate contextual information can not be retrieved from memory (and is not provided by some external source, such as a conversation), then comprehension is difficult. As for the concreteness effect in language comprehension, the context-availability model assumes that retrieving associated contextual information for abstract words is more difficult than for concrete items because abstract word representations presumably have weaker connections to contextual information than do representations for concrete words.

What follows is a series of simulations using the Hyperspace Analogue to Language (HAL) model (Burgess & Lund, 1997; Lund & Burgess, 1996) to examine cognitively-relevant differences between the representations for abstract and concrete words. One particular goal is to determine whether abstract and concrete word representations differ in the availability of associated contextual information stored in memory (as suggested by Schwanenflugel & Shoben). HAL is a context model that develops word meaning from global co-occurrence statistics extracted from human on-line language use. A ~320 million word corpus of Usenet text is the input stream from which HAL records weighted co-occurrence information for the 70,000 most frequent vocabulary items. The process of recording these co-occurrences allows for the formation of a co-occurrence matrix from which word vectors are derived. Mathematically, these vectors represent points in a high-dimensional space. The similarity between words corresponds inversely to inter-point distances, with the assumption that the more similar two words are, the closer their points in the high-dimensional space (see Burgess & Lund, 1997, for further discussion of the HAL methodology). Conceptually, each vector represents the entire learning history of a given word in the context of other words. We claim that these context vectors provide robust representations of a number of important aspects of word meaning (Burgess, Livesay, & Lund, 1998). HAL has been used to account for grammatical class distinctions and semantic affects on syntactic processing (Burgess & Lund, 1997), several semantic and associative priming effects (Lund & Burgess, 1996; Lund, Burgess, & Audet, 1996), the sort of semantic errors made by deep dyslexia patients (Buchanan, Burgess, Lund, 1996), and cerebral asymmetries in semantic memory processing (Burgess & Lund, 1998).

Experiment 1: Demonstrating Separation of Abstract and Concrete Words in the HAL Model

Semantic differentiation of a small set of dissociable abstract words (five emotional words, e.g., *love*, *sorrow*; five legal terms, e.g., *judge*, *law*) has been demonstrated before using HAL context vectors (Burgess & Lund, 1997). However, what remains unclear is whether these context vectors can be used to examine more systematically the proposed distinctions between abstract and concrete words. The goal of this first simulation is to provide new evidence that contextual representations extracted from HAL can be categorized along the concreteness dimension. In this simulation three larger sets of abstract and concrete words are subjected to multidimensional scaling (MDS) in order to show that the interword distances in the high-dimensional space can provide a basis for this categorization, thus providing evidence that HAL representations are relevant to the issues presented in the Introduction.

Method

Materials. Bleasdale (1987) presented a list of 80 concrete and 80 abstract words arranged in prime-target pairs (see his Appendix A). These words had been rated by undergraduate students for concreteness and imageability, with concrete primes and targets rated reliably higher than abstract primes and targets on both characteristics. Without regard to prime/target status, 159 of these words (the word "glutton" did not occur in the HAL matrix, and was not included in any simulations or analyses presented herein) were placed into a stimulus pool for possible inclusion into the following simulations and subsequent analyses. As the amount of information that can be displayed in a two-dimensional MDS in an interpretable manner is somewhat limited, we chose not to include all 159 words in a single MDS. Rather, three separate sets of 20 concrete and 20 abstract words were pseudo-randomly sampled (without replacement) from the larger pool, and these individual subsets were treated as representative of the entire pool of words. The only restriction placed on this sampling procedure was an attempt to avoid including highly related words in the same MDS as highly associated and/or similar words have a tendency to "pair off" and increase apparent dissociations between word categories within an MDS (e.g., *truth* and *false*, *son* and *daughter*, *democracy* and *government*) that might tend to exaggerate the examined effects.

Procedure. Global co-occurrence vectors were extracted from the HAL model for each set of 40 words. Each vector was treated as a set of coordinates in a high-dimensional Euclidean space, and for each set of words a MDS solution was computed. The hypothesis was that these word vectors, representing the interword distances for the chosen set of words, would operate as a similarity matrix (Lund & Burgess, 1996).

Results

Each similarity matrix was analyzed by a MDS algorithm that projects points from a high-dimensional space into a lower-dimensional space in a nonlinear fashion that attempts to preserve the distances between points. The lower-dimensional projection allows for the visualization of the spatial relationships between the co-occurrence vectors. The two-dimensional MDS solutions for the three subsets of words are shown in Figures 1a-1c. To further simplify interpretation of the information present in the MDS solutions, concrete and abstract words are represented by Cs or As, respectively.

Visual inspection of Figures 1a-1c suggests that concrete and abstract words were differentiated in the MDS solutions. In each case the abstract words appear to occupy a space separate from the concrete words. However, given the nature of the MDS procedure, in which an extreme reduction of dimensionality occurs when projecting data from a high-dimensional space down to only two dimensions, it was necessary to perform appropriate inferential statistics. In order to determine whether abstract words do exist in a separate (high-dimensional) space than concrete words an analysis of variance was performed separately for each set of 40 words which compared intragroup distances between these words with intergroup distances between these items. Distances between all combinations of word pairs within a group (i.e., concrete or abstract words) were calculated and compared to the distances between all combinations of word pairs between groups. In each of the three analyses concrete words were differentiated from abstract words, $F_{set1}(1, 778) = 5.59$, $p = .018$; $F_{set2}(1, 778) = 11.7$, $p = .0007$; $F_{set3}(1, 778) = 8.13$, $p = .01$. As well, abstract words were differentiated from concrete words, $F_{set1}(1, 778) = 143.7$, $p < .0001$; $F_{set2}(1, 778) = 74.23$, $p < .0001$; $F_{set3}(1, 778) = 116.1$, $p < .0001$.

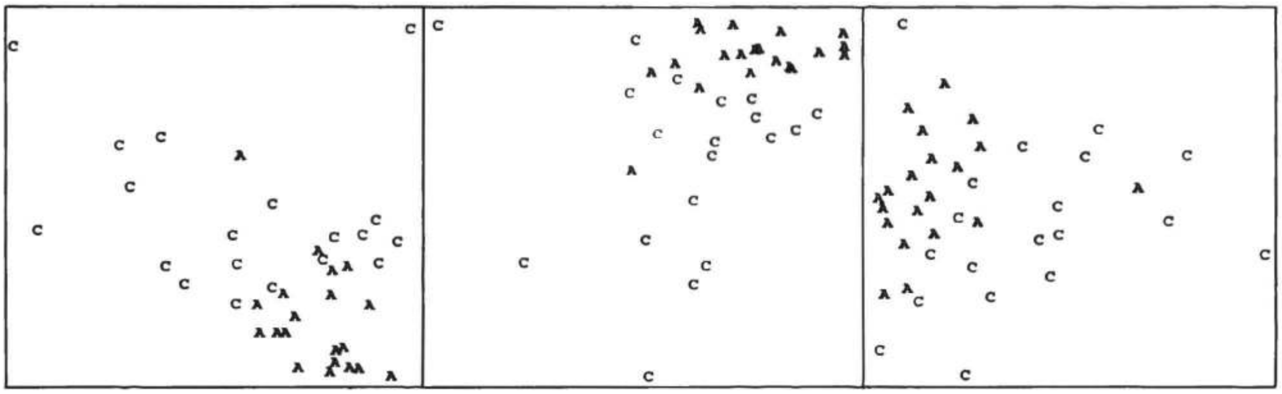

Figures 1a-1c: Two-dimensional multidimensional scaling solutions for word vectors for abstract (A) and concrete (C) words from Bleasdale (1987).

Discussion

As seen in Figures 1a-1c, it is clear that the information carried in HAL context vectors is sufficient to distinguish concrete from abstract words. This effect supports the range of theoretical and empirical evidence that suggests important differences in the processing and comprehension of these two word types. However, what remains unclear is the nature of the information contained in the HAL vectors that contributes to the present results. Moreover, it is also unclear what representational differences exist between abstract and concrete words for language users, differences that, presumably, play some important role in the language comprehension process. A number of possible explanations for the representational and processing differences between abstract and concrete words have been proposed (e.g., Paivio's dual-coding model, Schwanenflugel & Shoben's context-availability model; see Schwanenflugel, 1991 for a review of these proposals). Experiment 2 is an evaluation of one such theoretical model, the context-availability model, which attributes variations in word comprehension difficulty to differences in the retrieval of contextual information from memory. Given that HAL representations are derived from local and global contextual information it is expected that the context vectors extracted from this model will provide an avenue for evaluation of the context-availability model.

Experiment 2: Evaluating the Frequency and Context Diversity Hypotheses

Schwanenflugel and Shoben (1983; Schwanenflugel, 1991) suggest that concreteness effects in language comprehension can be explained with the context-availability model. As discussed in the Introduction, at the heart of this model is the idea that abstract words are more difficult to understand than concrete items because a person tends to have greater difficulty in retrieving associated contextual information from the "knowledge base" (i.e., the mental lexicon) for abstract words. In her discussion of this model Schwanenflugel (1991, p. 243) provided two distinct hypotheses regarding how abstract words might come to have weaker connections to contextual information.

Hypothesis 1: Abstract words occur so infrequently in language that representations for these words have relatively few opportunities to develop strong connections to contextual information stored in memory.

Hypothesis 2: Abstract words occur frequently, but in such a diversity of contexts that the opportunity to develop strong connections to one or a few particular contexts does not arise.

A corollary to the latter hypothesis is that although concrete words might occur in a relatively few number of contexts, these words are presumably well-grounded in one or a few contexts (perhaps due to strong connections to the visual environment, as Paivio would suggest), thus providing for secure links between concrete word representations and stored contextual information. As for the former hypothesis, it may be that more frequent words simply have more opportunities to establish stronger ties to stored contextual information, and concrete words may, in fact, be more frequent.

Materials. As Schwanenflugel did not provide a list of stimulus words in her articles, we chose to use the 159 words provided by Bleasdale (1987), along with an additional set of abstract and concrete words borrowed from Chiarello, Senehi and Nuding (1987) that were added to increase the scope of our investigation. Both Bleasdale and Chiarello et al. had subjects norm their stimuli for concreteness and imageability, demonstrating that their concrete words were rated reliably higher than abstract words on both characteristics. After removing duplicate words we were left with a list of 119 abstract and 129 concrete words.

Testing Hypothesis 1: Testing Schwanenflugel's first hypothesis regarding the possibly infrequent occurrence of abstract words in language was a simple matter of comparing raw frequency counts for the abstract and concrete items in the 320 million word HAL corpus (see Burgess & Livesay, 1998). Contrary to Hypothesis 1, abstract words were more frequent (228 occurrences per million words) than concrete words (136 occurrences per million), $t(246) = 2.84$,

p = .0049. This result shows that language users (specifically, English speakers) encounter abstract words much more frequently than concrete words, thus providing evidence that language users have more opportunities to develop strong connections between abstract words and associated contextual information, as compared to concrete words.

Testing Hypothesis 2: The disconfirmation of Hypothesis 1 provides some evidence for Schwanenflugel's second hypothesis: abstract words do occur more frequently than concrete words. However, this is not a direct indication that abstract words occur in a greater diversity of contexts. Examining this notion requires a measure that moves beyond raw word frequencies, and provides a metric of the different contexts in which a word was experienced by the HAL model. We believe that decomposing HAL's context vectors provides such a metric, in that each element in a word's context vector is a weighted co-occurrence count of that word with some other word in the text stream input to the HAL model. Considered more simply, each of the vector elements is a direct record of each of the contexts in which a target word occurred. If a target word never co-occurred with a particular word in the input stream--thus indicating that the target word was never experienced by the HAL model in that particular context--the vector element recording that potential co-occurrence would have a value of zero. On the other hand, any actual co-occurrences between the target word and some other word would result in the element representing that co-occurrence having a value greater than zero, thus representing the model's experience with that word in that context.

With the above consideration in mind, it is a simple matter to calculate the proportion of non-zero elements in a word's context vector (i.e., across the entire 70,000-element vector). This proportion--which we have labeled "context density"--directly represents the number of contexts in which a given word has been encountered by the HAL model across its entire learning history for that word. Analysis of the abstract and concrete words shows that the context density for the abstract words (16.8%) is greater than for concrete words (12.7%), $t(246) = 2.12$, $p = .035$, thus supporting Schwanenflugel's notion that abstract words occur in a greater diversity of contexts. In fact, the roughly 4% difference between the context densities for abstract and concrete words represents a difference of approximately 2,800 contexts--a rather substantial difference.

Discussion

Abstract words were shown to be more frequent than concrete words in a large sampling of human language use. This finding disconfirms Schwanenflugel's first hypothesis, and suggests that whatever differences between abstract and concrete word representations contribute to the concreteness effect in language comprehension, these differences are not attributable to a lack of experience with abstract words. Not only are abstract words more frequent than concrete words, abstract words also appear in a greater number of contexts. This latter finding supports Schwanenflugel's second hypothesis, and suggests that abstract word representations might have more diffuse connections to associated contextual information. To understand this conclusion, one might consider that a frequently occurring abstract word, such as "vacation," appears in such a variety of contexts that the representation for that word has few significant ties to any particular context. In contrast, a less frequent concrete word like "pyramid" is strongly associated with only a few distinct contexts (e.g., Egypt and pharaohs). We are aware that the *diversity* of contexts in which a word appears is not reflected only in the raw number of contexts for that word. Contextual diversity is also a function of the differences between the number of occurrences of a word in each of the different contexts in which that word was experienced. These different patterns of co-occurrence translate into different variances across the elements in HAL context vectors (Lund & Burgess, 1996). Put more simply, two words might occur in roughly the same number of contexts, but vary in the overall pattern of occurrences across these contexts, thus leading to quantitatively and qualitatively different representations.

Though the results from Experiment 2 do not fully address the veracity of the context-availability model as an adequate explanation of concreteness effects, these findings do support the basic assumption that abstract and concrete word representations differ on context availability (i.e., context density). We argue that the representational differences between abstract and concrete words are largely due to differences in how these words are used in language, and the context vectors extracted from the Hal model are transparently sensitive to such differences in word use. In fact, these results provide the first quantitative evidence that abstract and concrete words differ in the number of contexts in which these word types occur in natural language. What remains to be seen is whether differences in context density relate to the processing differences between abstract and concrete words seen in human studies. Experiment 3 involves a consideration of the concreteness effect as a function of the distinction between automatic and controlled processing.

Experiment 3: The Relationship Between Priming and Semantic Distance

As discussed in the Introduction, concreteness effects have been found using a number of experimental paradigms including both lower-level (e.g., lexical decision and naming) and higher-level tasks (e.g., sentence verification). However, few studies have examined concreteness effects along with an explicit consideration of the difference between lower-level and higher-level *processing*, in other words, the distinction between automatic and controlled processing. Two exceptions are the semantic priming studies presented by Bleasdale (1987) and Chiarello et al. (1987) in which the authors explicitly manipulated both word concreteness and the degree of automatic versus controlled processing. In a controlled priming paradigm using lexical decision (Exp. 1: 75% proportion of related prime-target pairs, and subject instructions designed to focus attention on prime word identities) Chiarello et al. showed greater priming for targets when preceded by a concrete prime, as compared to an abstract prime. This effect was not found in an automatic priming paradigm (Exp. 3: 25% proportion of related prime-target pairs; no instruction focusing attention on prime word identities) in which priming was equivalent for targets preceded by concrete or abstract primes. Bleasdale showed a similar pattern of results for both automatic and controlled priming. Bleasdale presented data from a naming (Exp. 1) and lexical decision task (Exp. 2), each involving controlled processing (i.e., a long stimulus-onset asynchrony between primes and targets, and a 50% proportion of related prime-target pairs), in which targets preceded by concrete primes showed greater priming than when preceded by abstract primes. In contrast, in a lexical decision task (Exp. 3) using an automatic priming procedure (i.e., a brief stimulus-onset asynchrony between primes and targets; a 50% proportion of related prime-target pairs) Bleasdale showed equivalent priming for targets following either a concrete or abstract prime.

In the following simulation we examine Bleasdale's prime-target pairs with the intent to show that the context vectors extracted from the HAL model mimic the initial bottom-up activation of semantic representations in memory. We have argued elsewhere (Burgess & Lund, in press) that the information carried by the context vectors best represents the sorts of information that are activated in memory early on in language processing (e.g., in word recognition). In fact, HAL context vectors are removed from any attentional or strategic effects that appear in controlled processing situations. Based on the results of Bleasdale and Chiarello et al. in their automatic processing paradigms, we expect to find equivalent context distance (i.e., semantic) priming with our vector representations when targets are paired with abstract or concrete primes.

Method

Materials. For our related condition we borrowed 79 prime-target pairs listed in Bleasdale's (1987) Appendix A (one pair was removed because the target did not occur in the HAL matrix). These stimuli consisted of 20 pairs with concrete primes and targets, 20 pairs with abstract primes and targets, 20 pairs with concrete primes and abstract targets, and 19 pairs with abstract primes and concrete targets. From these related pairs we generated 79 unrelated pairs by pseudo-randomly pairing each target with an unrelated prime. Though we have no way of directly comparing our unrelated pairs to those used by Bleasdale, when producing these pairs we did explicitly avoid producing unrelated pairs in which the prime and target shared any obvious relationship.

Procedure. Context vectors for all primes and targets were extracted from the HAL model. For each prime-target pair the Euclidean distance (in the high-dimensional context space represented by the HAL global co-occurrence matrix) between the two words was calculated. These distances are based on context vectors which have been normalized simply to provide a more easily interpretable set of values that map onto somewhat realistic word recognition reaction times. Analogous to the procedure used in human priming studies, in this simulation each target acted as its own control, with the context distance between a target and its related prime being compared to the distance between the target and its unrelated prime. A priming effect was then calculated for each word pair by subtracting the distance for the related pair from the distance for the unrelated pair.

Results

The comparison of primary interest was the difference in distance priming between targets that were paired with concrete primes compared to those paired with abstract primes. While an analysis of variance indicated an overall priming effect with related prime-target pairs displaying smaller context distances (603 distance units) than unrelated prime-target pairs (661 units), $F(1, 154) = 8.59$, $p = .004$, this effect did not interact with prime concreteness, $F < 1$. Thus, the distance priming seen for targets paired with abstract primes (53 units) did not differ from the distance priming found for targets paired with concrete primes (62 units).

Discussion

The pattern of results provided by Bleasdale and Chiarello et al. suggests that concreteness effects tend to appear when the experimental task promotes the use of higher-level, controlled processing of verbal stimuli, such as when subjects have an opportunity and the motivation to develop expectancies and response strategies in word recognition (Neely, 1991). The results of Experiment 3 using Bleasdale's stimuli replicate previous failures to find an effect of word concreteness in automatic processing paradigms. This supports our previous findings that HAL's context vectors provide a good match to empirical semantic priming results (Lund, Burgess, & Audet, 1995).

General Discussion

While it appears that processing differences do exist between abstract and concrete words, the nature of the representational differences between these word types that presumably underlie such processing differences is not well-understood. An evaluation of one proposed explanation of these representational and processing differences--the context-availability model--using word meaning representations provided by the HAL memory model suggests that important differences may exist between the diversity of contexts in which abstract and concrete words are experienced by language users. Our findings indicate that context diversity, as measured by context density in HAL's word representations (and perhaps context variance, as well), may have some utility in providing an explanation of concreteness effects in higher-level processing tasks. As for the issue of how context density might relate to the lack of concreteness effects in automatic processing paradigms, the predictions provided by the context-availability model are not motivated by differences in the time-course of word meaning activation, per se, but rather by the notion that contextual information associated with abstract words is simply more difficult to access and retrieve. We propose that the context distances between HAL context vectors are better predictors of early time-course effects in semantic memory activation than are variances in context diversity.

Acknowledgements

This research was supported by NSF Presidential Faculty Fellow award SBR-9453406 to Curt Burgess.

References

Balota, D. A., Ferraro, F. R., & Connor, L. T. (1991). On the early influence of meaning in word recognition: A review of the literature. In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 223-250). Hillsdale, NJ: Lawrence Erlbaum Associates.

Belmore, S. M., Yates, J. M., Bellack, D. R., Jones, S. N., & Rosenquist, S. E. (1982). Drawing inferences from concrete and abstract sentences. *Journal of Verbal Learning and Verbal Behavior, 21,* 338-351.

Bleasdale, F. A. (1987). Concreteness-dependent associative priming: Separate lexical organization for concrete and abstract words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 582-594.

Buchanan, L., Burgess, C., & Lund, K. (1996). Overcrowding in semantic neighborhoods: Modeling deep dyslexia. *Brain and Cognition, 32,* 111-114.

Burgess, C., & Livesay, K. (1998). The effect of corpus size in predicting reaction time in a basic word recognition task: Moving on from Kucera and Francis. *Behavior Research Methods, Instruments, and Computers, 30,* 272-277.

Burgess, C., & Lund, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12,* 1-34.

Burgess, C., & Lund, K. (1998). Modeling cerebral asymmetries in high-dimensional semantic space. In M. Beeman & C. Chiarello (Eds.), *Right hemisphere language comprehension: Perspectives from cognitive neuroscience.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Burgess, C., & Lund, K. (In press). The dynamics of meaning in memory. In E. Dietrich & A. Markman (Eds.) *Cognitive dynamics: Conceptual and representational change in humans and machines.*

Chiarello, C., Senehi, J., & Nuding, S. (1987). Semantic priming with abstract and concrete words: Differential asymmetry may be postlexical. *Brain and Language, 31,* 43-60.

Collins, A. M. & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407-28.

de Groot, A. M. B. (1989). Representational aspects of word imageability and word frequency assessed through word association. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 824-845.

Holmes, V. M., & Langford, J. (1976). Comprehension and recall of abstract and concrete sentences. *Journal of Verbal Learning and Verbal Behavior, 15,* 559-566.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers, 28,* 203-208.

Lund, K., Burgess, C., & Audet, C. (1996). Dissociating semantic and associative word relationships using high-dimensional semantic space. *Proceedings of the Cognitive Science Society* (pp. 603-608). Hillsdale, NJ: Lawrence Erlbaum Associates.

Neely, J. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G.W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Lawrence Erlbaum Associates.

Paivio, A. (1986). *Mental representations: A dual coding approach.* Oxford, England: Oxford University Press.

Ransdell, S. E., & Fischler, I. (1987). Memory in a monolingual mode: When are bilinguals at a disadvantage? *Journal of Memory and Language, 26,* 392-405.

Schwanenflugel, P. J. (1991). Why are abstract concepts hard to understand? In P. J. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 223-250). Hillsdale, NJ: Lawrence Erlbaum Associates.

Schwanenflugel, P. J., Akin, C., & Luh, W. (1992). Context availability and the recall of abstract and concrete words. *Memory and Cognition, 20,* 96-104.

Schwanenflugel, P. J., Harnishfeger, K. K., & Stowe, R. W. (1988). Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language, 27,* 499-520.

Schwanenflugel, P. J., & Shoben, E. J. (1983). Differential context effects in the comprehension of abstract and concrete materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 82-102.

Schwanenflugel, P. J., & Stowe, R. W. (1989). Context availability and the processing of abstract and concrete words. *Reading Research Quarterly, 24,* 114-126.