2013

Semantic Memory

Ken McRae

Michael Jones

SEMANTIC MEMORY

Ken McRae

University of Western Ontario

Michael Jones

Indiana University

Ken McRae
Department of Psychology
Social Science Centre
University of Western Ontario
London, ON, Canada N6A 5C2

Email: kenm@uwo.ca
Phone: (519) 661-2111(x84688)
Fax: (519) 661-3961

Michael Jones
Department of Psychology
1101 E. 10th St.
Indiana University
Bloomington, IN, 47405

Email: jonesmn@indiana.edu
Phone: (812) 856-1490
Fax: (812) 855-4691

1. INTRODUCTION

Concepts and meaning are fundamental components of nearly all aspects of human cognition. We use this knowledge every day to recognize entities and objects in our environment, anticipate how they will behave and interact with each other, use them to perform functions, to generate expectancies for situations, and to interpret language. This general knowledge of meaning falls within the realm of *semantic memory*. For many years, semantic memory was viewed as an amodal, modular memory store for factual information about concepts, distinct from episodic memory (our memory for specific instances of personal experience). However, researchers now interpret semantic memory more broadly to refer to general world knowledge, entangled in experience, and dependent on culture. Furthermore, there is now considerable evidence suggesting that semantic memory is grounded in the sensory modalities, is distributed across brain regions, and depends on episodic memories at least in terms of learning, with the possibility that there is no definite line between episodic and semantic memory. In this chapter, we review contemporary research in semantic memory. We limit our discussion to lexical semantics (the meaning of individual words), with particular focus on recent findings and trends, formal computational models, neural organization, and future directions.

1.1. Classic View of Semantic Memory

Tulving (1972) viewed memory as a system of independent modules. Long-term memory was subdivided into declarative (facts) and procedural (skills) components. Declarative memory was further divided into semantic memory and episodic memory, with a clear distinction between them. Tulving characterized semantic memory as amodal. In an amodal view, when one thinks of an apple, the information retrieved from semantic memory is independent of the sensory modalities used to perceive an apple. Although semantic memory contains factual information about an apple's color and taste, this information is dissociated from the sensory systems used to actually see or taste.

Early neuropsychological evidence supported Tulving's (1972) view. For example, amnesic patients showed dissociations between episodic and semantic memory tasks (Squire, 1988); their impairment seemed to have little effect on semantic memory despite profound episodic deficiencies, bolstering the modularity claim. Research on 'schema abstraction' tasks found that the decay of episodic and semantic memory followed different profiles (Posner & Keele, 1968). Although memory for episodes is stronger than for category prototypes immediately after training, episodic memory decays much faster than semantic memory (or at least, instances decay faster than do abstract prototypes).

Tulving's (1972) characterization of semantic memory as an amodal, modular system separate from episodic and procedural memory provided a useful foundation to study and understand human semantic representations. In retrospect, however, it may have actually stifled research in semantics by imposing a rigid framework that is unlikely to be correct. Recent research with improved experimental, computational, and neuroimaging techniques clearly contradicts the classic view. Semantic memory is now viewed more broadly as a part of an integrated memory system, grounded in the sensory, perceptual, and motor systems, and is distributed across key brain regions.

2. GROUNDING SEMANTIC MEMORY

Tulving's classic view of semantic memory as an amodal symbolic store has been challenged by contemporary research. There is a growing body of behavioral and neuroimaging research demonstrating that when humans access word meaning, they automatically activate sensorimotor information used to perceive and act on the real-world objects and relations to which a word refers. In theories of *grounded cognition*, the meaning of a word is grounded in the sensorimotor systems (Barsalou, 1999; see Pecher & Zwann, 2005, for a review). Hence, when one thinks of an apple, knowledge regarding motoric grasping, chewing, sights, sounds, and tastes used to encode episodic experiences of an apple are reinstated via sensorimotor simulation. Thus, a grounded simulation refers

to context-specific re-activations that incorporate the important aspects of episodic experience into a current representation. In this sense, simulations are guided and only partial (Barsalou, 2008). This approach challenges amodal views, and makes a clear link between episodic experience and semantic memory.

A wealth of recent behavioral evidence supports the grounded simulation approach to semantics. For example, response latencies for images and feature names are faster when they have visual properties congruent with context (Solomon & Barsalou, 2001; Zwann, Stanfield, & Yaxley, 2002). Similarly, having participants perform particular motions (e.g., grasping) facilitates the comprehension of sentences describing actions involving these motions (Klatzky et al., 1989), and prime-target pairs sharing motor-manipulation features (e.g., *typewriter-piano*) are responded to more quickly than pairs that do not (Myun, Blumstein, & Sedivy, 2006). Zwann and Madden (2005) review numerous studies suggesting that the mental representations activated during comprehension also include information about object features, temporal and spatial perspective, and spatial iconicity. Barsalou (2008) and Pecher, Boot, and Van Danzig (2011) contain surveys of the recent literature attesting to the importance of situation models, simulation (perceptual, motor, and affective), and gesture in language comprehension and abstract concepts.

3. SEMANTIC ORGANIZATION IN THE BRAIN

Semantic memory research was for many years dominated by cognitive psychologists who generally were not concerned with neural organization. In cognitive neuropsychology, there is a history of studies investigating patients with semantic deficits (Warrington & Shallice, 1984). However, for a number of years, this line of research was divorced from semantic memory research using normal adult participants. With the advent of neuroimaging techniques, fMRI in particular, research on the neural organization of semantic memory blossomed.

Researchers have long known that brain regions responsible for perception tend to be specialized for specific sensory modalities. Given that perception is distributed across specialized neural regions, one possibility is that conceptual representations are organized in a similar fashion. For the past 40 years, Paivio (1971) has advocated a form of modality-specific representations in his dual-coding theory. Furthermore, studies of patients with category-specific semantic deficits have been used as a basis for arguing for multimodal representations for the past 25 years or so. In early work, Warrington and McCarthy (1987) put forward their sensory/functional theory to account for patterns of category specific impairments of knowledge in patients with focal brain damage. The basic assumption is that living things depend primarily on visual knowledge, whereas although visual knowledge is also important for nonliving things, knowledge of an object's function is primary. Building on Allport (1985), recent research has used analyses of large scale feature production norms to extend the sensory-functional theory to other senses and types of knowledge, and move beyond the binary living-nonliving distinction (Cree & McRae, 2003). There do remain some accounts of category-specific semantic deficits that are amodal (Caramazza & Shelton, 1998, Tyler & Moss, 2001), but even these researchers have begun to find support for theories in which knowledge is tied to modality-specific brain areas (Mahon & Caramazza, 2003; Raposo, Moss, Stamatakis, & Tyler, 2009).

The behavioral and neuropsychological evidence in favor of grounded semantics is corroborated by recent neuroimaging studies supporting a distributed multimodal system. A few researchers have used evoked response potentials to investigate this issue (Sitnikova, West, Kuperberg, & Holcomb, 2006), but the vast majority of studies have used fMRI. For example, Goldberg, Perfetti, and Schneider (2006) tied together previously reported neuroimaging evidence supporting modally bound tactile, colour, auditory, and gustatory representations. They found that sensory brain areas for each modality are recruited during a feature verification task using linguistic stimuli (e.g., *banana-yellow*). The same pattern emerges in single word processing. Hauk, Johnsrude, and Pulvermüller (2004) showed that reading action words correlates with activation in somatotopically corresponding areas of the motor

cortex (*lick* activates tongue regions while *kick* activates foot regions), indicating that word meaning is modally distributed across brain regions. Furthermore, within brain regions that encode modality-specific, possibly feature-based representations, some studies suggest a category-based organization (Chao, Haxby, & Martin, 1999). Finally, some studies have shown that semantic representations are located just anterior to primary perceptual or motor areas, whereas others have found evidence for activation of primary areas (see Thompson-Schill, 2003). In summary, there is a large amount of converging evidence supporting a distributed multimodal semantic system (for thorough reviews, see Binder, 2009; Martin, 2007).

Perhaps one the most important remaining issues concerns the fact that people's concepts are not experienced as a jumble of features, disjointed across space and time, but instead are experienced as coherent unified wholes. Multimodal feature-based theories therefore need to include a solution to the binding problem, specifying how representational elements are integrated into conceptual wholes, both within and between modalities. One solution involves temporal synchrony of neuronal firing rates (von der Malsburg, 1999). Semantic representations may be integrated by coincidental firing rates of distributed neural populations. However, the most frequently invoked solution is based on the idea of a convergence zone, which can be considered as a set of processing units that encode coincidental activity among multiple input units (Damasio, 1989). In connectionist models, a convergence zone may be thought of as a hidden layer (Rogers et al., 2004). Because they encode time-locked activation patterns, an important property of convergence zones is that they transform their input, rather than just repeat signals. In this way, successive convergence zones build more complex or abstract representations. Current theories of multimodal semantic representations incorporate either single convergence zones, as in Patterson, Nestor, and Rogers' (2007) anterior temporal lobe hub theory, or a hierarchy of convergence zones encoding information over successively more complex configurations of modalities (Simmons & Barsalou, 2003). At the moment, it is unclear which of these hypotheses is correct.

In summary, recent research supports the idea that semantic representations are grounded across modality-specific brain regions. Researchers are working toward fleshing out details of precisely what these regions encode, the degree to which sub-regions are specific to types of concepts, and how semantic representations are experienced as unified wholes. Furthermore, the vast majority of research has been conducted on concrete concepts, so research on other concepts, such as verbs or abstract concepts, will play a key role over the next few years.

4. EVENT-BASED SEMANTIC REPRESENTATIONS

Another way in which the semantic-episodic distinction has been blurred in recent years concerns research on event-based knowledge in semantic memory. People's knowledge of common everyday events includes actions that are part of those events, and common primary participants or components, such as agents (the people doing the action), patients (the people or objects upon which the action is performed), instruments involved in actions, and locations at which various events take part. Furthermore, people have knowledge of temporal aspects of events. This generalized event knowledge is learned through our experience with everyday events, watching television and movies, and reading and hearing about what people have done, what they are doing, and what they are going to do.

Language provides many cues into event knowledge. For example, verbs like *travel* or *cook* denote events and actions, some nouns like *breakfast* refer to events, and other nouns refer to entities or objects that typically play a role in specific situations, such as *waitress*, *customer*, *fork*, or *cafeteria*. A number of studies have shown that such event knowledge is computed rapidly from single words. These experiments have tended to use a priming paradigm with a short stimulus onset asynchony (SOA: the time between the onset of the prime and the onset of the target), which is viewed as providing a window into the organization of semantic memory.

Moss, Ostrin, Tyler, and Marslen-Wilson (1995) showed priming effects based on instrument relations (such as *broom-floor*) and what they called script relations, in which the primes were a mixture of events and locations (*hospital-doctor* and *war-army*). Subsequent studies have shown that verbs prime their typical agents (*arresting-cop*), patients (*serving-customer*), and instruments (*stirred-spoon*), but not locations (*skated-arena*; Ferretti, McRae, & Hatherell, 2001). Furthermore, typical agents, patients, instruments, and locations prime verbs (McRae, Hare, Elman, & Ferretti, 2005). In addition, Hare, Jones, Thomson, Kelly, and McRae (2009) showed that event nouns prime the types of people and things commonly found at those events (*sale-shopper*, *breakfast-eggs*), location nouns prime entities and objects typically found at those locations (*stable-horse*, *sandbox-shovel*), and instrument nouns prime the types of things on which they typically are used (*key-door*) but not the people who typically use them (*hose-gardener*, although priming was found in the other direction). Hare et al. used a corpus-based model, BEAGLE (Jones, Kintsch, & Mewhort, 2006) to simulate their results.

Chwilla and Kolk (2005) showed that people can integrate words rapidly to construct situations, thus producing priming. They presented two words simultaneously that were unrelated except when considered in the context of some broader event (*director bribe*), and demonstrated priming of a third word (*dismissal*) related to the situation. Chwilla and Kolk's results depend on conceptually integrating both primes with the target, thus speaking to rapid activation of knowledge of situations. In addition, Khalkhali, Wammes, and McRae (2011) found that relatedness decision latencies were shorter when three events were presented in the order corresponding to their usual real-world sequence (*marinate-grill-chew*) than when the order of the first two events was reversed (*grill-marinate-chew*), suggesting that such temporal information is encoded in semantic memory.

An interesting consequence of these studies is that they move toward a stronger tie between semantic memory and sentence comprehension. For example, a number of the studies used thematic roles of verbs as the basis for testing relations, thus making direct contact with a key construct in sentence processing research. Along this same line, Jones and Love (2007) provide a point of contact between sentence processing and how people learn lexical concepts. Participants studied sentences such as *The polar bear chases the seal* and *The German shepherd chases the cat*. In a test phase, similarity ratings for entities and objects participating in common relational systems increased. The increase was largest for objects playing the same role within a relation (e.g., the chaser), but also was present for those playing different roles in the same relation (e.g., the chaser or the chasee role in the chase relation), and this happened regardless of whether they participated in the same sentence/event.

In summary, recent studies have investigated people's episodic-based knowledge of common generalized events. These studies show that semantic memory is organized so that this knowledge is computed and used rapidly, and they demonstrate direct links between episodic and semantic memory.

5. SEMANTIC AND ASSOCIATIVE RELATIONS

There are longstanding issues in semantic memory research regarding associative versus semantic relations. Association has a long history in psychology and philosophy, and normative word association has often been used to explain performance in semantic memory experiments (Nelson, McEvoy, & Dennis, 2000; Roediger, Watson, McDermott, & Gallo, 2001). Bower (2000) defined associations as "sensations that are experienced contiguously in time and/or space. The memory that sensory quality or event A was experienced together with, or immediately preceding, sensory quality or event B is recorded in the memory bank as an association from idea *a* to idea *b*." (p. 3). In 1965, Deese stated that "almost all the basic propositions of current association theory derive from the sequential nature of events in human experience" (p. 1). More recently, Moss et al. (1995) claimed that associations between words are "built up through repeated co-occurrence of the two word forms." (p. 864). In general, the consensus seems to be that contiguity is key to forming a link between two concepts.

In contrast, association in cognitive psychology almost invariably is defined in terms of its operationalization. In a word association task, a participant hears or reads a stimulus word and produces the first word that comes to mind (Nelson, McEvoy, & Schreiber, 1998). Thus, two words are associated if one is produced as a response to the other. There exist significant discrepancies between the definition of association and its operationalization. Association is learning-based whereas word association is production-based. Association is based on sensory information whereas word association is linguistically-based. Association is based on contiguity, whereas word associations are virtually always meaningful.

The construct of semantic relatedness was for a long time limited to exemplars from the same category, or featurally-similar concepts (as in *cow-horse*; Lupker, 1984; Frenck-Mestre & Bueno, 1999). Recently, however, researchers have investigated a much wider array of relations. The event-based relations discussed in Section 4 are examples. In addition, researchers have been studying what are often called thematic relations (see Estes, Golonka, & Jones, 2011, for a recent review). These include, for example, *cow-milk*, where a cow produces milk, or *wind-erosion*, where wind causes erosion.

A thorny issue concerns delineating between the influences of semantic and associative relatedness. Lucas (2000) concluded from a meta-analysis of priming experiments that "pure" semantic priming (in the absence of word association) exists, whereas there is no evidence for association-based priming in the absence of semantic relatedness. In contrast, Hutchinson (2003) reviewed individual studies and concluded that both semantic and associative relatedness produce priming. One possibility is that it may not be fruitful to distinguish between associative and semantic relations because word associations are best understood in terms of semantic relations (Anisfeld & Knapp, 1968; Brainerd et al., 2008). In some views, the word association task unambiguously taps associative connections between words/concepts in people's semantic memory. In contrast, word association can be considered an open-ended task on which performance is driven almost exclusively by types of semantic relations. Researchers who have classified word associates according to their semantic relations have shown that almost all stimulus-response pairs, with the exception of rhymes, have clear semantic relations (Guida & Lenci, 2007). Furthermore, Brainerd et al. found that a number of semantic variables correlate with word association strength.

This is likely the primary reason why it has been so difficult to distinguish empirically between associative and semantic relations. In studies of associative priming, the items are a mixture of semantic relations, such as *hammer-nail* or *engine-car*. McNamara (2005) stated the issue clearly: "Having devoted a fair amount of time perusing free-association norms, I challenge anyone to find two highly associated words that are not semantically related in some plausible way. Under this view, the distinction between purely semantically and associatively related words is an artificial categorization of an underlying continuum." (p. 86). Furthermore, in studies of pure semantic relatedness priming, items that appear in word association norms are excluded. However, it does not appear to make sense to argue that items in these studies are not associated in the general sense. For example, Hare et al. (2009) analyzed subsets of stimuli not associated according to word association norms, showing priming in the absence of association. This logic appears at first glance to be valid because concepts such as *sale* and *shopper* are not associated according to Nelson et al.'s (1998) norms. However, shoppers are found at sales, and the entire point of a sale is to attract shoppers. So, these concepts definitely are associated in the general sense, even though forward and backward association statistics indicate that they are not.

The line between association and semantics has now been questioned in a number of areas of research. McRae, Khalkhali, and Hare (in press) discuss this issue with respect to research using the Deese-Roediger-McDermott false memory paradigm, picture-word facilitation and interference, the development of associative learning through adolescence, and semantic priming. Although association is a critical aspect of learning, one possibility is that virtually all retained associations are meaningful and thus can be understood in terms of semantic relations. On the other hand, given longstanding views

on the primacy of association-based links in memory (as indexed by normative word association), this debate is likely to continue.

6. ABSTRACT CONCEPTS

The structure and content of abstract concepts such as *lucky*, *advise*, and *boredom* have been studied to a much lesser extent than have concrete concepts, and thus are not nearly as well understood. In general, there is little consensus regarding how abstract concepts are represented and computed. The lack of obvious physical referents in the world for abstract concepts makes theorizing, model building, and experiment quite difficult, but also an important and intriguing issue. We use the phrase "obvious physical referents" in the previous sentence because many abstract concepts are at least partly experienced by the senses, or have internal states that correspond to them. For example, we have all experienced boredom, we have internal thoughts and emotions that are tied to the meaning of *boredom*, and we can visually recognize boredom in other people.

The most influential theory has been Paivio's (1971; 2007) dual-coding theory, in which the processing of lexical concepts involves the activation of functionally independent but interconnected verbal and nonverbal representational systems. The verbal system consists of associatively interconnected linguistically-based units, whereas the nonverbal system consists of spatially-organized representations of objects and events that can be experienced as mental images. Activation spreads within and between systems. Concrete concepts are represented in both systems, whereas abstract concepts are represented in the verbal system only. Dual-coding theory has been used to explain differences between concrete and abstract words in memory tasks, lexical decision, EEG, and fMRI experiments.

Dual-coding theory is contrasted frequently with context-availability theory, in which the major difference between abstract and concrete concepts is that abstract words and sentences are more difficult to comprehend because it is challenging to access relevant world knowledge contextual information when comprehending abstract materials (Schwanenflugel & Shoben, 1983). At present, however, dual-coding theory has received much more support.

The vast majority of experiments on abstract concepts compare performance on concrete versus abstract words, either in isolation or in sentence contexts. A consistent finding is that memory is better for concrete concepts (Paivio, 2007). A number of studies have also found shorter lexical decision latencies to concrete than to abstract words in isolation (Schwanenflugel & Shoben, 1983), and a larger N400 to concrete words (Kounios & Holcomb, 1994). Some patients have been reported with better performance on concrete concepts (Coltheart, Patterson, & Marshall, 1980). However, a frustrating aspect of this research is that, although the memory results are stable, some studies show shorter lexical decision latencies to abstract words (Kousta et al., 2011), and some patients perform better on abstract concepts (Breedin, Saffran, & Coslett, 1994). In addition, there is no compelling explanation for the N400 results. Finally, the fMRI literature on concrete versus abstract concepts has produced highly variable results (Grossman et al., 2002; Kiehl et al., 1999; Wise et al., 2000).

There are at least two reasons for the inconsistency in results. First, some differences may be task related because the manner in which people process words influences the form of concrete-abstract differences. Second, there may be important differences among item sets across studies. Typically, researchers select concrete and abstract words using concreteness and/or imageability ratings. However, the categories of concrete and abstract concepts are large, and selecting small subsets from these large classes has presumably led to inconsistent results. To deal with this issue, researchers have begun to classify abstract words on further dimensions, such as emotional valence (Kousta et al., 2011).

More recent theories of the structure and content of abstract concepts have emerged. In Barsalou's (1999) perceptual simulation theory, abstract and concrete concepts can be simulated from prior experience. One issue involves the type of simulations that might be key to abstract concepts that do not, at least at first glance, have sensory-motor correspondences. Barsalou and Weimer-Hastings

(2005) focused on situations as the key to abstract concepts. Concepts such as *lucky*, *advise*, and *boredom* are tied both to situations in which people have learned the meaning of those concepts, and to internally generated cognitive and emotional states. At present, however, little research has been conducted to flesh out these ideas.

One other prominent theory of abstract concepts is conceptual metaphor or image schema theory (Lakoff, 1987). In this view, abstract concepts are mapped to sensory-motor grounded image schemas. For example, studies suggest that the abstract concept of time is grounded in our knowledge of space (Casasanto & Boroditsky, 2008). At the moment, however, the notion of a conceptual metaphor or image schema is inconsistent among theorists (Pecher et al., 2011).

A promising avenue for studying abstract concepts comes from corpus-based distributional models. One advantage of corpus-based models is that they provide representations for all types of words using the same computational mechanisms. As described in Section 7.3, these models can also be combined with other approaches to form hybrids.

In summary, understanding the organization and content of abstract concepts is a major challenge for all current theories of semantic memory. Addressing the relevant issues will require a deeper appreciation of the similarities and differences among types of abstract concepts, how abstract concepts depend differentially on sensory, motor, and internally-generated cognition and emotional information, and the degree to which they are tied to situations or contexts in which they are important.

7. COMPUTATIONAL MODELS OF SEMANTIC MEMORY

There is a fuzzy boundary in the literature between models of semantic processing and semantic representation. We define the former to be models of how learned semantic structure is used in tasks, and the latter to be models that specify a mechanism by which semantic memory is formed from experience. We first review models of semantic processing, and then models of learning semantic structure from experience (primarily corpus-based models). However, we acknowledge that this distinction between the two 'levels' of models is an oversimplification: How we learn new semantic information depends on the current contents of semantic memory, and semantic structure and process influence each other when explaining behavioral data (Johns & Jones, 2010).

7.1. Models of Semantic Processing

Connectionist networks have been used to provide insight into how word meaning is represented and computed, and to simulate numerous empirical semantic memory phenomena. In these models, concepts are typically represented as distributed patterns of activity across sets of representational units that often represent features (<has 4 legs>), but not necessarily nameable ones. Units are organized into layers, and are connected by weighted connections. These connections control processing, and their weights are established using a learning algorithm.

The impact of connectionist models has been at least four-fold. First, due to distributed representations, they naturally encode concept similarity in terms of shared units, and thus simulate similarity-based phenomena (Cree et al., 1999; Masson, 1995; Plaut & Booth, 2000). Second, because they learn statistical regularities between and within patterns, they have led researchers to focus on the distributional statistics underlying semantic representations and computations (Tyler & Moss, 2001; McRae, de Sa, & Seidenberg, 1997). Third, because many connectionist models settle into representations over time (e.g., attractor networks), they can be used to simulate response latencies, and provide insight into the temporal dynamics of computing word meaning (Masson, 1995). Fourth, one can train a model and then damage it in various ways, thus simulating brain-damaged patients (Hinton & Shallice, 1991; Plaut & Shallice, 1993; Rogers et al., 2004). Finally, all of these properties of connectionist models are interrelated.

Semantic processing unfolds over time. When we read or hear a word, components of meaning become active at different rates over the first several hundred milliseconds. Attractor networks, in

which units update their states continuously based on both their prior states and input from other units, are well suited to simulate this process. Because priming has played such a large role in semantic memory research, a number of researchers have simulated it. Given that similar concepts have overlapping distributed representations, connectionist networks have been successful at simulating priming between featurally-similar concepts such as *eagle* and *hawk*, providing insight into factors such as correlations among semantic features and the degree of similarity between concepts (Cree et al., 1999; McRae et al., 1997). Furthermore, researchers have simulated contiguity-based (associative) priming, and individual differences in priming (Masson, 1995; Plaut & Booth, 2000).

One way in which distributional statistics underlying semantic representations have been studied is the feature verification task, in which participants judge whether a feature such as <has an engine> is reasonably true of a concept such as *van*. These studies and accompanying simulations have highlighted the role of correlational structure. That is, some features tend to occur with others across basic-level concepts, such as <has feathers> and <has a beak>, and there is a continuum of the feature correlational strength. Studies such as McRae et al. (1997) and Randall et al. (2004) show that connectionist models predict influences of feature correlations that are observed in human data. Furthermore, the degree to which features are distinctive (the inverse of the number of concepts in which a feature occurs) plays a privileged role in semantic computations in both people and connectionist simulations (Cree et al., 2006; Randall et al., 2004). Distributional statistics such as these are bases for theories such as the conceptual structure account (Tyler & Moss, 2001), and are also strongly implicated in understanding data from category-specific deficits and semantic dementia patients (Rogers et al., 2004; Tyler & Moss). Finally, they may form the basis for understanding how superordinate categories such as *clothing* and *fruit* are learned and computed (O'Connor, Cree, & McRae, 2009).

Much of the research on simulating neurally-impaired adults has drawn on work by Hinton and Shallice (1993) and Plaut and Shallice (1993). A nice example is Rogers and colleagues' work in which they provide detailed accounts of semantic dementia patients (Rogers et al., 2004; Rogers & McClelland, 2004). Rogers et al. damaged a trained attractor network, and then simulated patient performance in a number of tasks. For example, they showed that loss of knowledge followed a specific-to-general trajectory because of the nature in which regularities across visual and verbal patterns are stored in their model's hidden units. Features that were shared and correlated across numerous concepts tended to be represented in larger and more neighboring regions of semantic space than were highly distinctive features. Thus, distinctive features were more likely to be influenced by damage, so the model showed a tendency to lose its ability to discriminate among similar concepts early in the course of semantic dementia.

Finally, Rogers and McClelland (2004) present a large set of arguments and simulations in which, among other things, they provide connectionist accounts of several phenomena that have been highlighted in knowledge-based theories of concepts. The issues are too complex and numerous to do them justice in a short paragraph, but their book is highly recommended.

7.2 Models of Semantic Representation

Classic models of semantic structure assumed that meaning was represented either as a hierarchical network of interconnected nodes (Collins and Quillian, 1969) or as arrays of binary features (Smith, Shoben, & Rips, 1974). A major limitation of both these early models is that neither specifies how their representations are learned. Instead, their representations must be hand coded by the researcher or collected from adult participants.

More recent *distributional models* specify cognitive mechanisms for constructing semantic representations from statistical experience with text corpora. In general, these models are all based on the *distributional hypothesis* (Harris, 1970): Words that appear in similar linguistic contexts are likely to have related meanings. For example, *apple* may frequently co-occur with *seed, worm,* and *core.* As a

result, the model can infer that these words are related. In addition, the model can induce that *apple* is similar to *peach* even if the two never directly co-occur, because they occur around the same types of words. In contrast, *apple* and *staple* rarely appear in the same or similar contexts.

There are a large number of distributional models (for reviews, see Bullinaria & Levy 2007; Riordan & Jones, 2011). To simplify discussion, we classify them into three families based on their learning mechanism: 1) latent inference models, 2) passive co-occurrence models, and 3) retrieval-based models. For an in-depth review of new advances in distributional modeling, we refer readers to the recent pair of special issues of *Topics in Cognitive Science*, (2011).

7.2.1. Latent Inference Models. This family of models reverse-engineers the cognitive variables responsible for how words co-occur across linguistic contexts. The process is similar to other types of latent inference in psychology. For example, personality psychologists commonly administer structured questionnaires, constructing items to tap hypothetical psychological constructs. Singular value decomposition (SVD) is applied to the pattern of responses over questionnaire items to infer the latent psychological variables responsible for the cross-item response patterns. Latent inference models of semantic memory work in an analogous way, but they apply this decomposition to the pattern of word co-occurrences over documents in a corpus.

The best-known latent inference model is *Latent Semantic Analysis* (LSA; Landauer & Dumais, 1997). LSA begins with a word-by-document frequency matrix from a text corpus. Each word is weighted relative to its entropy over documents; 'promiscuous' words appearing in many contexts are dampened more than are 'monogamous' words that appear more faithfully in particular contexts. Finally, the matrix is factored using SVD, and only the components with the largest eigenvalues are retained (typically 300-400). These are the latent semantic components that best explain how words co-occur over documents, similar to the way that the psychological constructs of introversion and extroversion might explain response patterns over hundreds of questionnaire items. With this reduced representation, each word in the corpus is represented as a pattern over latent variables. In the reduced space, indirect relationships emerge—even though two words may never have directly co-occurred in a document (e.g., two synonyms), they can have similar patterns.

Landauer and Dumais (1997) suggested that the human brain performs some data reduction operation similar to SVD on contextual experience to construct semantic representations. However, they were careful not to claim that what the brain does is *exactly* SVD on a perfectly remembered item-by-episode representation of experience. Whether or not LSA is a plausible model of human semantic representation (for criticisms, see Glenberg & Robertson, 2000; Perfetti, 1998), it has been remarkably successful at accounting for data ranging from human performance on synonymy tests (Landauer & Dumais, 1997) to metaphor comprehension (Kintsch, 2000). LSA set the stage for future distributional models to better study the specific mechanisms that might produce a reduced semantic space. In addition, the model made a clear formal link between semantic memory structure and episodic experience.

More recently, Griffiths, Steyvers, and Tenenbaum's (2007) *Topic* model extended LSA in a Bayesian framework, specifying a generative mechanism by which latent semantic variables could produce the pattern of word co-occurrences across documents. The Topic model operates on the same initial data representation as LSA—it assumes that we experience words over discrete episodic contexts (operationalized as documents in a corpus). However, it specifies a cognitive inference process based on probabilistic reasoning to discover word meaning. To novice users of semantic models, the computational machinery of the Topic model can be daunting. However, the theoretical underpinning of the model is simple and elegant, and is based on the same idea posited for how children infer unseen causes for observable events. Consider an analogy: Given a set of co-occurring symptoms, a dermatologist must infer the unseen disease or diseases that led to the observed symptoms. Over many instances of the same co-occurring symptoms, she can infer the likelihood that they result from a

common cause. The topic model works in an analogous manner, but on a much larger scale of inference and with mixtures of causal variables. Given that certain words tend to co-occur in contexts and this pattern is consistent over many contexts, the model infers the likely latent "topics" that are responsible for the co-occurrence patterns, where the document is a probabilistic mixture of topics. A word's meaning is a probability distribution over possible topics, where a topic is a probability distribution over words (just as a disease would be a probability distribution over symptoms, and a symptom is a probability distribution over possible diseases that produced it).

This results in two key distinctions from LSA. First, the Topic model is generative in that it defines a process by which documents could be constructed from mixtures of mental variables. Second, a word's representation is a probability distribution rather than a point in semantic space. This allows the Topic model to represent multiple meanings of ambiguous words, whereas in LSA, ambiguity is collapsed to a single point. The Topic model is able to account for free association data, sense disambiguation, word-prediction, and discourse effects that are problematic for LSA (Griffiths et al., 2007).

7.2.2. Passive Co-Occurrence Models.
Passive co-occurrence models posit simple Hebbian type accumulation mechanisms that give rise to sophisticated semantic representations. Hence, these models tend not to need a full word-by-document matrix, but gradually develop semantic structure by simple co-occurrence "counting" as a text corpus is continuously experienced.

The first passive co-occurrence model was the *hyperspace analogue to language* model (HAL; Lund & Burgess, 1996). HAL slides an *n*-word window across a text corpus, and counts the co-occurrence frequency of words within the window (where frequency is inversely proportionate to distance between words in the window). A word's semantic representation is a vector of distance-weighted co-occurrence values to all other words in the corpus. Hence, a word is defined relative to other words in HAL. Comparing the vectors for two words yields their semantic similarity, producing both direct and indirect semantic relations as in LSA. HAL has accounted for a range of semantic priming phenomena (Lund & Burgess, 1996). Modern variants of HAL have improved the model to produce better fits to human data (Rhode, Gonnerman, & Plaut, 2009), and a HAL-like model was used by Mitchell et al. (2008) to predicted fMRI activation patterns associated with the meanings of concrete nouns.

A second type of passive co-occurrence model uses the accumulation of random vectors as a mechanism for semantic abstraction (Jones, Kintsch, & Mewhort, 2006; Kanerva, 2009). For example, in the BEAGLE model (Jones & Mewhort, 2007) words are initially represented by random patterns of arbitrary dimensionality. Hence, before any episodic experience, the representation for *apple* is no more similar to *peach* than it is to *staple*. As text is experienced, each word's memory pattern is updated as the sum of the random initial patterns representing the words with which it co-occurs. Thus, *apple, peach,* and *core* move closer to one another in semantic space as text is experienced, while *staple* moves away (but closer to *paper, pencil,* etc.). Random accumulation can be considered as semantic abstraction from the coincidental co-occurrence of (initially random) brain states representing the words in the episodic context. Because of the arbitrary nature of the features, BEAGLE can simultaneously learn about the positional information of words in the context similar to HAL-type models (Jones & Mewhort, 2007 use convolution to encode order information). Hence, a word's representation becomes a pattern of arbitrary "features" that reflects its history of co-occurrence with, and position relative to, other words in contexts. BEAGLE simulates a number of phenomena including semantic priming, typicality, and semantic constraints in sentence completions.

7.2.3. Retrieval-Based Models.
Rather than assuming that humans store semantic representations, retrieval-based models construct meaning as part of episodic memory retrieval. Retrieval-based models are similar to exemplar-based theories of categorization (Nosofsky, 1986) and multiple-trace theories

of global memory (Hintzman, 1986). Just as Hintzman's Minerva 2 model demonstrated that schema abstraction is simulated by a model containing only episodic traces, Kwantes' (2005) constructed semantics model demonstrates that semantic phenomena are possible without requiring a semantic memory. In this model, memory *is* the episodic word-by-context matrix. When a word is read or heard, its semantic representation is constructed as an average of other words in memory, weighted by their contextual similarity to the target. Although semantic abstraction differs radically from LSA, similar representations are produced. Dennis (2005) has used a similar approach, accounting for an impressive array of semantic phenomena.

7.3. Integrating Perceptual Information into Distributional Models
Distributional models have been criticized as psychologically implausible because they learn from only linguistic information and do not contain information about sensorimotor perception contrary to grounded cognition (for a review, see de Vega, Glenberg, & Graesser, 2008). Hence, representations in distributional models are not a replacement for feature norms. Feature-based representations contain a great deal of sensorimotor features of words that cannot be learned from purely linguistic input, and both types of information are core to human semantic representation (Louwerse, 2008). Riordan and Jones (2011) recently compared a variety of feature-based and distributional models on semantic clustering tasks. Their results demonstrated that whereas there is information about word meaning redundantly coded in both feature norms and linguistic data, each has its own unique variance and the two information sources serve as complimentary cues to meaning.

Research using recurrent networks trained on child-directed speech corpora has found that pretraining a network with features related to children's sensorimotor experience produced significantly better word learning when subsequently trained on linguistic data (Howell, Jankowicz, & Becker, 2005). Durda, Buchanan, and Caron (2009) trained a feedforward network to associate LSA-type semantic vectors with their corresponding activation of features from McRae et al.'s (2005) norms. Given the semantic representation for *dog,* the model attempts to activate the output feature <has fur> and inhibit the output feature for <made of metal>. After training, the network was able to infer the correct pattern of perceptual features for words that were not used in training because of their linguistic similarity to words that were used in training.

A recent flurry of models using the Bayesian Topic model framework has also explored parallel learning of linguistic and featural information (Andrews, Vigliocco, & Vinson, 2009; Baroni, Murphy, Barba, & Poesio, 2010; Steyvers, 2009). Given a word-by-document representation of a text corpus and a word-by-feature representation of feature production norms, the models learn a word's meaning by simultaneously considering inference across documents and features. This enables learning from joint distributional information: If the model learns from the feature norms that *sparrows* have beaks, and from linguistic experience that *sparrows* and *mockingbirds* are distributionally similar, it infer that *mockingbirds* also have beaks, despite having no feature vector for *mockingbird*. Integration of linguistic and sensorimotor information allows the model to better fit human semantic data than a model trained with only one source (Andrews et al., 2009). This information integration is not unique to Bayesian models, but can also be accomplished within passive co-occurrence models (Jones & Recchia, 2010; Vigliocco, Vinson, Lewis, & Garrett, 2004) and retrieval-based models (Johns & Jones, in press).

8. SUMMARY
Over the past 25 years or so, semantic memory research has blossomed for a number of reasons, all of which are equally important. The generation of intriguing patient data and theories of the organization of semantic memory that resulted from them, and the acceptance of this patient research into what some might call mainstream cognitive psychology, was an important step. Furthermore, connectionist models of semantic processing enabled implementations of meaning-based computations, generating

new ideas, experiments, and simulations. The advent of neuroimaging methods allowed researchers to study semantic processing in the brain, to integrate neurally-based theories with those resulting from implemented models as well as normal adult and patient data, and to generate novel theories of semantic representation and processing. In addition, theories of grounded cognition added excitement and paved the way for a large number of novel experiments designed to test them. Finally, corpus-based models of meaning have provided new ways to think about semantic representations, and a plethora of new ideas for designing experiments, and techniques for simulating human performance. The present high level of enthusiasm surrounding the study of semantic memory should continue as researchers refine, compare, and integrate theories, and test predictions that result from those theoretical endeavors. We hope that we have communicated some of this excitement to the reader.

References

Allport, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current Perspectives in Dysphasia* (pp. 207-244). Edinburgh: Churchill Livingstone.

Andrews, M., Vigliocco, G. & Vinson, D. P. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review, 116*, 463-498.

Anisfeld, M., & Knapp, M. (1968). Association, synonymity, and directionality in false recognition. *Journal of Experimental Psychology, 77*, 171-179.

Baroni, M., Murphy, B., Barbu, E., & Poesio, M. (2010). Strudel: A corpus-based semantic model based on properties and types. *Cognitive Science 34,* 222-254.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral Brain Science, 22*, 577-660.

Barsalou, L. W. (2008). Grounding symbolic operations in the brain's modal systems. In G. R. Sermin & E. R. Smith (Eds.), *Embodied grounding: Social, cognitive, affective, and neuroscientific approaches* (pp. 9-42). New York, NY: Cambridge University Press.

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking* (pp. 129–163). Cambridge: Cambridge University Press.

Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A review of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*, 2767-2796.

Bower, G. H. (2000). A brief history of memory research. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford Handbook of Memory* (pp. 3-32). New York: Oxford University Press .

Brainerd, C. J., Yang, Y., Reyna, V. F., Howe, M. L., & Mills, B. A. (2008). Semantic processing in "associative" false memory. *Psychonomic Bulletin & Review, 15*, 1035-1053.

Breedin, S. D., Saffran, E. M., & Coslett, H. B. (1994). Reversal of the concreteness effect in a patient with semantic dementia. *Cognitive Neuropsychology*, 11, 617–660.

Bullinaria, J. A. & Levy, J.P. (2007). Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods, 39*, 510-526.

Caramazza, A., & Shelton, J.R. (1998). Domain specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience, 10*, 1-34.

Casasanto, D., & Boroditsky, L. (2008). Time in the mind: Using space to think about time. *Cognition, 106,* 579–593.

Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. Nature Neuroscience, 2, 913-919.

Chwilla, D. J., & Kolk, H. H. J. (2005). Accessing world knowledge: Evidence from N400 and reaction time priming. *Cognitive Brain Research, 25*, 589-606.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240-247.

Coltheart, M., Patterson, K., & Marshall, J. (1980). *Deep dyslexia*. London: Routledge.

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.

Cree, G. S., McNorgan, C., & McRae, K. (2006). Distinctive features hold a privileged status in the computation of word meaning: Implications for theories of semantic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 643-658.

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

Damasio, A. R. (1989). The brain binds entitites and events by multiregional activation from convergence zones. *Neural Computation, 1*, 123-132.

De Vega, M., Glenberg, A. M., & Graesser, A. C. (2008). *Symbols and embodiment: debates on meaning and cognition.* Oxford, UK: Oxford University Press.

Deese, J. (1965). *The structure of associations in language and thought.* Baltimore, MD: Johns Hopkins Press.

Dennis, S. (2005). A memory-based theory of verbal cognition. *Cognitive Science, 29*, 145-193.

Durda, K., Buchanan, L., & Caron, R. (2009). Grounding co-occurrence: Identifying features in a lexical co-occurrence model of semantic memory. *Behavior Research Methods, 41*, 1210-1223.

Estes, Z., Golonka, S., & Jones, L. L. (2011). Thematic thinking: The apprehension and consequences of thematic relations. In B. Ross (Ed.), The Psychology of Learning & Motivation, Vol. 54 (pp. 249-278). Burlington: Academic Press.

Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory & Language, 44*, 516-547.

Frenck-Mestre, C., & Bueno, S. (1999). Semantic features and semantic categories: Differences in the rapid activation of the lexicon. *Brain and Language, 68*, 199-204.

Glenberg, A., & Robertson, D. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language, 43*, 379-401.

Goldberg, R. F., Perfetti, C. A., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *The Journal of Neuroscience, 26*, 4917-4921.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review, 114*, 211-244.

Grossman, M., Koenig, P., DeVita, C., Glosser, G., Alsop, D., Detre, J., & Gee, J. (2002). The neural basis for category-specific knowledge: An fMRI study. *Neuroimage*, *15*, 936-948.

Guida, A., & Lenci, A. (2007). Semantic properties of word associations to Italian verbs. *Italian Journal of Linguistics*, 19, 293-326.

Hare, M., Jones, M. N., Thomson, C., Kelly, S., & McRae, K. (2009). Activating event knowledge. *Cognition, 111*, 151-167.

Harris, Z. (1970). Distributional structure. In *Papers in Structural and Transformational Linguistics* (pp. 775-794). Dordrecht, Holland: D. Reidel Publishing Company.

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron, 41,* 301-307.

Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review, 98,* 74-95.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411-428.

Howell, S., Jankowicz, D., & Becker, S. (2005). A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language, 53*, 258-276.

Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review, 10*, 785-813.

Johns, B. T., & Jones, M. N. (2010). Evaluating the random representation assumption of lexical semantics in cognitive models. *Psychonomic Bulletin & Review, 17,* 662-672.

Johns, B. T., & Jones, M. N. (in press). Perceptual inference from global lexical similarity. *Topics in Cognitive Science.*

Jones, M. & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology*, *55*, 196-231.

Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language, 55*, 534-552.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review, 114*, 1-37.

Jones, M. N., & Recchia, G. L. (2010). You can't wear a coat rack: A binding framework to avoid illusory feature migrations in perceptually grounded semantic models. *Proceedings of the 32nd Annual Cognitive Science Society*, (pp. 877-882). Austin TX: CSS.

Khalkhali, S., Wammes, J., & McRae, K. (2011). Integrating words denoting typical sequences of events. *Canadian Journal of Experimental Psychology*.

Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representations with high-dimensional random vectors. *Cognitive Computation, 1*, 139-159.

Kiehl, K. A., Liddle, P. F., Smith, A. M., Mendrek, A., Forster, B. B., & Hare, R. D. (1999). Neural pathways involved in the processing of concrete and abstract words. *Human Brain Mapping, 7*, 225-233.

Kintsch, W. (2000). Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review, 7,* 257-266.

Klatzky, R. L., Pellegrino, J. W., McCloskey, B. P., & Doherty, S. (1989). Can you squeeze a tomato? The role of motor representations in semantic sensibility judgments. *Journal of Memory & Language, 28,* 56-77.

Kounios, J., & Holcomb, P. J. (1994). Concreteness effects in semantic processing: ERP evidence supporting dual-coding theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 804-823.

Kousta, S., Vigliocco, G., Del Campo, E., Vinson, D. P., & Andrews, M. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, *140*, 14-34.

Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin & Review, 12,* 703-710.

Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: Chicago University Press.

Landauer, T. K., & Dumais, S. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211-240.

Louwerse, M. M. (2007). Symbolic or embodied representations: A case for symbol interdependency. In T. K. Landauer, D. S. Macnamara, S. Dennis & W. Kintsch (Eds.), *Handbook of Latent Semantic Analysis* (pp. 107-120). Mahwah, NJ: Erlbaum.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review, 7*, 618-630.

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

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

Mahon, B. Z. & Caramazza, A. (2003). Constraining questions about the organization and representation of conceptual knowledge. *Cognitive Neuropsychology*, 20, 433-450.

Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology, 58,* 25-45.

Masson, M. E. J., (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 3-23.

McNamara, T. (2005). *Semantic Priming*. New York, NY: Psychology Press.

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

McRae, K., Hare, M., Elman, J. L., & Ferretti, T. R. (2005). A basis for generating expectancies for verbs from nouns. *Memory & Cognition, 33*, 1174-1184.

McRae, K., Khalkhali, S., & Hare, M. (in press). Semantic and associative relations: Examining a tenuous dichotomy. In V. F. Reyna, S. Chapman, M. Dougherty, & J. Confrey (Eds.), *The adolescent brain: Learning, reasoning, and decision making*. Washington, DC: APA.

Mitchell, T. M., Shinkanerva, S. V., Carlson, A., Chang, K., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *Science, 320,* 1191-1195.

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.

Nelson, D. L., McEvoy, C. L., & Dennis, S. (2000). What is free association and what does it measure? *Memory & Cognition, 28*, 887-899.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. <http://web.usf.edu/FreeAssociation/>

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.

O'Connor, C. M., Cree, G. S., & McRae, K. (2009). Conceptual hierarchies in a flat attractor network: Dynamics of learning and computations. *Cognitive Science*, *33*, 665-708.

Paivio, A. U. (1971). *Imagery and Verbal Processes*. Holt, Rinehart, and Winston: New York.

Paivio, A. U. (2007). *Mind and its Evolution: A Dual Coding Theoretical Approach*. Erlbaum: Mahwah, NJ.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, *8*, 976-987.

Pecher, D., Boot, I., & Van Dantsig, S. (2011). Abstract concepts: Sensory-motor grounding, conceptual metaphors, and beyond. In B. Ross (Ed.), The Psychology of Learning & Motivation, Vol. 54 (pp. 217-248). Burlington: Academic Press.

Pecher, D. & Zwaan, R. A. (2005). *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press.

Perfetti, C. (1998). The limits of co-occurrence: Tools and theories in language research. *Discourse Processes, 25*, 363-377.

Plaut, D. C., & Booth, J. R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review, 107*, 786-823.

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

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychlogy, 77,* 353-363.

Randall, B., Moss, H. E., Rodd, J., Greer, M., & Tyler, L. K. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. *Journal of Experimental Psychology: Language, Memory, and Cognition, 30,* 393-406.

Raposo, A., Moss, H. E., Stamatakis, E. A., & Tyler, L. K. (2009). Modulation of the motor and premotor cortices by actions, action words, and action sentences. *Neuropsychologia*, *47*, 388-396.

Rohde, D. L. T., Gonnerman, L., & Plaut, D. C**.** (2009). An improved model of semantic similarity based on lexical co-occurrence. *Cognitive Science.*

Riordan, B., & Jones, M. N. (2011). Redundancy in linguistic and perceptual experience: Comparing distributional and feature-based models of semantic representation. *Topics in Cognitive Science, 3, 303-345.*

Roediger, H. L., III, Watson, J. M., McDermott, K. B., & Gallo, D. A. (2001). Factors that determine false recall: A multiple regression analysis. *Psychonomic Bulletin & Review, 8*, 385-407.

Rogers, T. T. , Lambon Ralph, M. A, Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., and Patterson, K. (2004). The structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review*, 111, 205-235.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A parallel distributed processing approach*. MIT Press.

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

Simmons, K. W., & Barsalou, L. (2003). The similarity-in-topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology, 20,* 451-486.

Sitnikova T., West, W. C., Kuperberg, G. R., & Holcomb, P. J. (2006). The neural organization of semantic memory: Electrophysiological activity suggests feature-based segregation. *Biological Psychology, 71*, 326-340.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81*, 214-241.

Solomon, K. O., & Barsalou, L. W. (2001). Representing properties locally. *Cognitive Psychology, 43,* 129-169.

Steyvers, M. (2009). Combining feature norms and text data with topic models. *Acta Psychologica, 133*, 234-243.

Squire, L. R. (1988). Episodic memory, semantic memory, and amnesia. *Hippocampus, 8,* 205-211.

Thompson-Schill, S. (2003). Neuroimaging studies of semantic memory: Inferring how from where. *Neuropsychologia*, *41*, 280-292.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381-403). New York: Academic Press .

Tyler, L. K., & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences*, *5*, 244-252.

Vigliocco, G., Vinson, D. P., Lewis, W., & Garrett, M. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology, 48*(4), 422-488.

von der Malsburg, C. (1999). The what and why of binding: The modeler's perspective. *Neuron, 24*, 95 – 104.

Warrington, E. K., & McCarthy, R. A. (1987). Categories of knowledge. Further fractionation and an attempted integration. *Brain*, *110*, 1273-1296.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829-854.

Wise, R. J. S., Howard, D., Mummery, C. J., Fletcher, P., Leff, A., Buchel, C., & Scott, S. K. (2000). Noun imageability and the temporal lobes. *Neuropsychologia*, *38*, 985-994.

Zwann, R. A., Stanfield, R. A., & Yaxley, R. H. (2002). Language comprehenders mentally represent the shape of objects. *Psychological Science, 13,* 168-171.

Zwaan, R. A., & Madden, C. J. (2005). Embodied sentence comprehension. In Pecher, D. & Zwaan, R. A. (Eds.) *Grounding cognition: The role of perception and action in memory, language, and thinking*. (pp. 224-245). Cambridge, UK: Cambridge University Press.