

## 2.28 Semantic Memory

D. A. Balota and J. H. Coane, Washington University, St. Louis, MO, USA

© 2008 Elsevier Ltd. All rights reserved.

2.28.1	Nature of the Representation	512
2.28.2	Network Approaches	512
2.28.3	Feature Analytic Approaches	516
2.28.4	Concept Learning and Categorization	518
2.28.5	Grounding Semantics	520
2.28.5.1	Grounding Semantics in Analyses of Large-Scale Databases	520
2.28.5.2	Grounding Semantics in Perceptual Motor Systems	521
2.28.6	Measuring Semantic Representations and Processes: Insights from Semantic Priming Studies	522
2.28.7	The Interplay Between Semantics and Episodic Memory	525
2.28.8	Representation and Distinctions: Evidence from Neuropsychology	527
2.28.8.1	Category-Specific Deficits	527
2.28.8.2	Semantic Dementia	528
2.28.9	Neuroimaging	529
2.28.10	Development and Bilingualism	531
2.28.11	Closing Comments	531
References		531

Semantic memory entails the enormous storehouse of knowledge that all humans have available. To begin with, simply consider the information stored about the words of one's native language. Each of us has approximately 50,000 words stored in our mental dictionary. With each entry, we also have many different dimensions available. For example, with the word 'dog' we have stored information about how to spell it, how to pronounce it, its grammatical category, and the fact that the object the word refers to typically has four legs, is furry, is a common pet, and likes to chase cats (sometimes cars, squirrels, and other rodents), along with additional sensory information about how it feels when petted, the sound produced when it barks, the visual appearance of different types of dogs, emotional responses from past experiences, and much, much more. Of course, our knowledge about words is only the tip of the iceberg of the knowledge we have available. For example, people (both private and public) are a particularly rich source of knowledge. Consider how easy it is to quickly and efficiently retrieve detailed characteristics about John F. Kennedy, Marilyn Monroe, Bill Clinton, a sibling, parent, child, and so on. Indeed, our semantic, encyclopedic knowledge about the world appears limitless.

One concern reflected by the examples above is that semantic memory seems to be all inclusive. In this light, it is useful to contrast it with other forms of memory, and this is precisely what [Tulving \(1972\)](#) did in his classic paper distinguishing semantic and episodic memory. According to Tulving, semantic memory "is a mental thesaurus, organized knowledge a person possesses about words and other verbal symbols, their meaning and referents, about relations among them, and about rules, formulas, and algorithms for the manipulation of these symbols, concepts and relations" (1972, p. 386). In contrast, episodic memory refers to a person's memory for specific events that were personally experienced and remembered. So, the memory for the experience of having breakfast yesterday (e.g., where one was seated, how one felt, the taste of the food, who one was with) would fall under the umbrella of episodic memory, but the fact that eggs, cereal, and toast are typical breakfast foods reflects semantic knowledge. However, as we shall see, there is some controversy regarding where episodic memory ends and semantic memory begins. Indeed, we would argue that semantic memory penetrates all forms of memory, even sensory and working memory ([Sperling, 1960](#); [Tulving and Pearlstone, 1966](#); [Baddeley, 2000](#)), because tasks that are assumed to

tap into these other types of memory often are influenced by semantic memory.

So, what is indeed unique about semantic memory, and how has this area of research contributed to our understanding of learning and memory in general? One issue that researchers in this area have seriously tackled is the nature of representation, which touches on issues that have long plagued the philosophy of knowledge or epistemology. Specifically, what does it mean to know something? What does it mean to represent the meaning of a word, such as DOG? Is it simply some central tendency of past experiences with DOGS that one has been exposed to (i.e., a prototype DOG), or is there a limited list of primitive semantic features that humans use to capture the meaning of DOG, along with many other concepts and objects? Is the knowledge stored in an abstracted, amodal form that is accessible via different routes or systems, or is all knowledge grounded in specific modalities? For example, the meaning of DOG might be represented by traces laid down by the perceptual motor systems that were engaged when we have interacted with DOGS in the past.

In this chapter, we attempt to provide an overview of the major areas of research addressing the nature of semantic memory, emphasizing the major themes that have historically been at the center of research. Clearly, given the space limitations, the goal here is to introduce the reader to these issues and provide references to more detailed reviews. The vast majority of this work emphasizes behavioral approaches to the study of semantics, but we also touch upon contributions from neuropsychology, neuroimaging, and computational linguistics that have been quite informative recently. We focus on the following major historical developments: (1) the nature of the representation, (2) conceptual development and learning, (3) insights from and limitations of semantic priming studies, (4) interplay between semantic and episodic memory tasks, and (5) cognitive neuroscience constraints afforded by comparisons of different patient populations and recent evidence from neuroimaging studies. For further discussion of this latter area, the interested reader should see Chapter 2.29.

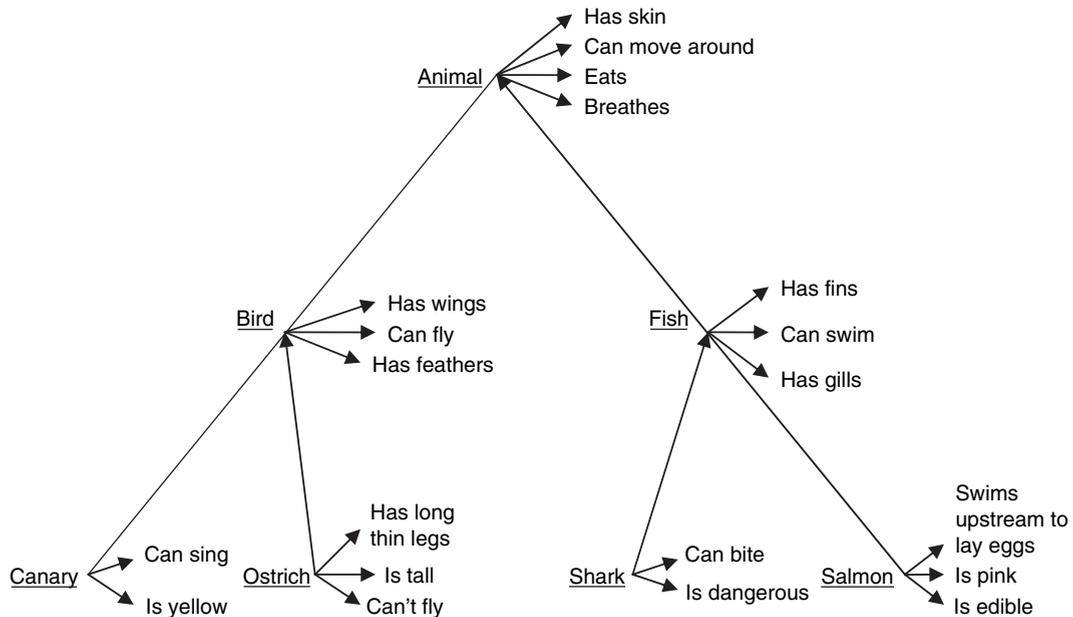
### **2.28.1 Nature of the Representation**

Although the question of how one represents knowledge has been around since the time of Aristotle, it is clear that cognitive scientists are still actively pursuing this issue. One approach to representation

is that we abstract from experience a prototypical meaning of a concept, and these ideal representations are interconnected to other related representations within a rich network of semantic knowledge. This is the network approach. Another approach is that there is a set of primitive features that we use to define the meaning of words. The meanings of different words and concepts reflect different combinations of these primitive features. This is a feature-based approach. Historically, the distinction between these two approaches has been central to research addressing the nature of semantic memory.

### **2.28.2 Network Approaches**

One of the first landmark studies of knowledge representation came from computer science and was based on the important dissertation of [A. M. Quillian \(1968\)](#) developed a model of knowledge representation called the Teachable Language Comprehender. A goal of this model was to formulate a working program that allowed efficient access to an enormous amount of information while minimizing redundancy of information in the network. Quillian adopted a hierarchically organized network, a portion of which is displayed in [Figure 1](#). As shown, there are two important aspects to the network: nodes and pathways. The nodes in this network are intended to directly represent a concept in semantic memory, so for example, the word BIRD has a node that represents BIRDNES. These nodes are interconnected in this network via labeled pathways, which are either 'isa' directional pathways or property pathways. Specifically, one can verify that BIRDS are indeed ANIMALS by finding an isa pathway between BIRDS and ANIMALS. Likewise, one could verify that 'A ROBIN BREATHES' by finding the isa pathway between robin and bird, and between bird and animal, and then accessing the property pathway leading to BREATHES from ANIMALS. In this sense, the model was quite economical, because most properties were stored only at the highest level in the network in which most of the lower exemplars included that property. For example, BREATHES would only be stored at the ANIMAL level, and not at the BIRD or CANARY level, thereby minimizing redundancy (and memory storage) in the network. Quillian also recognized that some features may not apply to all exemplars below that level in the network (e.g., ostriches are birds, and birds fly), so in these cases, one needed to include a



**Figure 1** Hierarchically arranged network. Taken from Collins A and Quillian MR (1969) Retrieval time from semantic memory. *J. Verb. Learn. Verb. Behav.* 8: 240–247.

special property for these concepts (such as CAN'T FLY attached to OSTRICHES).

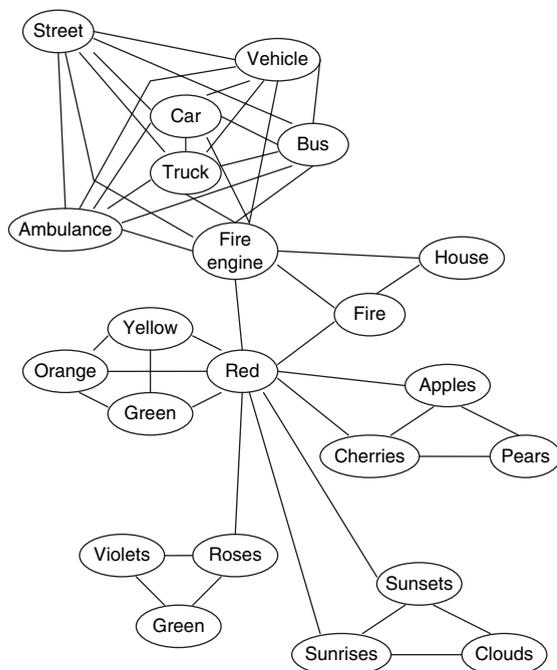
The economy of the network displayed in **Figure 1** does not come without some cost. Specifically, why would one search so deeply in a network to verify a property of a given concept, that is, why would one have to go all the way to the ANIMAL concept to verify that 'CANARIES BREATHE'? It seems more plausible that we would have the property BREATHES directly stored with the CANARY node. Of course, Quillian was not initially interested in how well his network might capture performance in humans, because his goal was to develop a computer model that would be able to verify a multitude of questions about natural categories, within the constraints of precious computer memory available at the time.

Fortunately for cognitive psychologists, Quillian began a collaborative effort with A. Collins to test whether the network model developed by Quillian could indeed predict human performance on a sentence verification task, that is, the speed to verify such sentences as 'A CANARY IS A BIRD'. Remarkably, the Collins and Quillian (1969) study provided evidence that appeared to be highly supportive of the hierarchically organized network structure that Quillian independently developed in artificial intelligence. Specifically, human performance was nicely predicted by how many 'isa' and

'property' pathways one needed to traverse to verify a sentence. The notion is that there was a spreading activation retrieval mechanism that spread across links within the network, and the more links traversed the slower the retrieval time. So, the original evidence appeared to support the counterintuitive prediction that subjects indeed needed to go through the 'CANARY IS A BIRD' link and then the 'BIRD IS AN ANIMAL' link to verify that 'CANARIES BREATHE', because this is where BREATHES is located in the network.

The power of network theory to economically represent the relations among a large amount of information and the confirmation of the counterintuitive predictions via the sentence verification studies by Collins and Quillian (1969) clearly encouraged researchers to investigate the potential of these networks. However, it soon became clear that the initial hierarchically arranged network structure had some limitations. For example, the model encountered some difficulties handling the systematic differences in false reaction times, that is, the finding that correct 'false' responses to 'BUTTERFLIES ARE BIRDS' are slower than responses to 'SPIDERS ARE BIRDS.' Importantly, there was also clear evidence of typicality effects within categories. Specifically, categories have graded structure, that is, some examples of BIRD, such as ROBINS, appear to be better examples than other BIRDS, such as OSTRICHES.

There were numerous attempts to preserve the basic network structure of Collins and Quillian (1969), and indeed, some general models of cognitive performance still include aspects of such network structure. Collins and Loftus (1975) took a major step forward when they developed a network that was not forced into a hierarchical framework. This is displayed in Figure 2. As shown, these networks are basically unstructured, with pathways between concepts that are related and the strength of the relationship being reflected by the length of the pathways. Collins and Loftus further proposed that the links between nodes could be dependent on semantic similarity (e.g., items from the same category, such as DOG and CAT, would be linked), or the links could emerge from lexical level factors, such as cooccurrence in the language. Thus, DOG and CAT would be linked because these two items often occur in similar contexts. Because the strength of spreading activation is a function of the distance the activation traversed, typicality effects can be nicely captured in this framework by the length of the pathways. Of course, one might be concerned that such networks are not sufficiently constrained by independent evidence (i.e., if one is slow the pathway must be long). Nevertheless, such networks have been implemented



**Figure 2** Semantic network. From Collins AM and Loftus EF (1975) A spreading-activation theory of semantic processing. *Psychol. Rev.* 82: 407–428.

to capture knowledge representation in both semantic and episodic domains (see Anderson, 2000).

More recently, there has been a resurgent interest in a type of network theory. Interestingly, these developments are again driven from fields outside of psychology such as physics (see Albert and Barabasi, 2002) and biology (Jeong et al., 2000). This approach is very principled in nature in that it uses large existing databases to establish the connections across nodes within a network and then uses graph analytic approaches to provide quantitative estimates that capture the nature of the networks. In this light, researchers are not arbitrarily constructing the networks but are allowing the known relations among items within the network to specify the structure of the network. This approach has been used to quantify such diverse networks as the power grid of the Western United States and the neural network of the worm, *C. elegans* (Watts and Strogatz, 1998). Once one has the network established for a given domain (i.e., providing connections between nodes), one can then quantify various characteristics of the network, such as the number of nodes, the number of pathways, the average number of pathways from a node, and the average distance between two nodes. Moreover, there are more sophisticated measures available such as the clustering coefficient, which reflects the probability that two neighbors of a randomly selected node will be neighbors of each other. In this sense, these parameters quantify the characteristics of the targeted network. For example, when looking at such parameters, Watts and Strogatz (1998) found that naturally occurring networks have a substantially higher clustering coefficient and relatively short average distances between nodes compared with randomly generated networks that have the same number of nodes and average connectivity between nodes. This general characteristic of networks is called ‘small world’ structure. These high clustering coefficients may reflect ‘hubs’ of connectivity and allow one to access vast amounts of information by retrieving information along the hubs. In popular parlance, such hubs may allow one to capture the six degrees of separation between any two individuals that Milgram (1967) proposed and that has been popularized by the game “six degrees of separation with Kevin Bacon”.

What do worms, power grids, and parlor games have to do with semantic memory? Steyvers and Tenenbaum (2005) used three large databases reflecting the meaning of words to construct networks of semantic memory. These included free-association

norms (Nelson et al., 1998), WordNet (Miller, 1990), and Roget's Thesaurus (1911). For example, if subjects are likely to produce a word in response to another word in the Nelson et al. free-association norms, then a connection between the two nodes was established in the network. Interestingly, Steyvers and Tenenbaum found that these semantic networks exhibited the same small world structure as other naturally occurring networks; specifically, high-clustering coefficients and a relatively small average path distance between two nodes. As shown in Figure 3, if one moves along the hub of highly interconnected nodes, an enormous amount of information becomes readily available via traversing a small number of links.

Of course, it is not a coincidence that naturally occurring networks have small world structure. The seductive conclusion here is that knowledge representation has some systematic similarities across domains. Indeed, Steyvers and Tenenbaum (2005) and others have suggested that such structure reflects central principles in development and representation of knowledge. Specifically, Steyvers and Tenenbaum argue that as the network grows, new nodes are predisposed to attach to existing nodes in a probabilistic manner. It is indeed quite rare that a new meaning of a word is acquired without it being some variation of a preexisting meaning (see Carey,

1978). Hence, across time, nodes that are added to the network will be preferentially attached to existing nodes. This will give rise to a high degree of local clustering, which is a signature of small world network structure. We return to the issue of how concepts develop in a later section.

It is noteworthy that Steyvers and Tenenbaum (2005) have also provided empirical support from their network analyses. For example, they have found that word frequency, or the degree to which a word is encountered in language, and age of acquisition, defined as the average age at which a child learns a given word, effects in naming and lexical decision performance naturally fall from this perspective. Naming and lexical decision are two of the most commonly used word recognition tasks used in research investigating the nature and structure of semantic memory. In naming (or speeded pronunciation), a participant is asked to read a presented stimulus aloud as quickly as possible, whereas in lexical decision, he or she is asked to indicate whether a letter string is a real word or a pseudoword (i.e., a string of letters that does not correspond to the spelling of a real word). In both tasks, the primary dependent measure is response latency. The general assumption is that the speed required to access the pronunciation of a word or to recognize a string of letters reflects processes involved in accessing stored

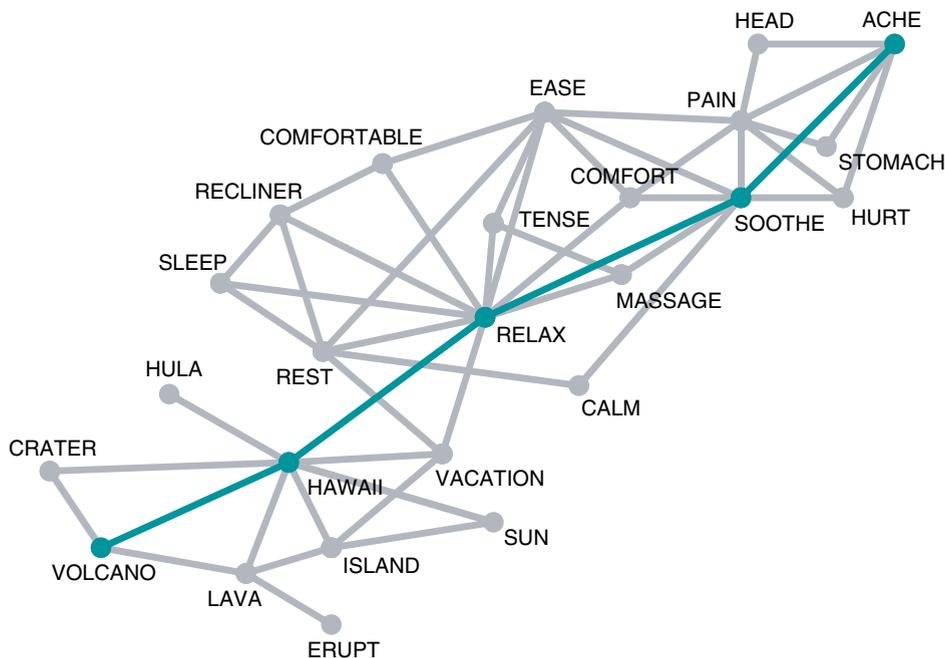


Figure 3 Segment of small world semantic network. Courtesy of Marc Steyvers.

knowledge about that word. Interestingly, Steyvers and Tenenbaum found a reliable negative correlation between number of connections to a node (semantic centrality) in these networks and response latency, precisely as one might predict, after correlated variables such as word frequency and age of acquisition have been partialled out (also see Balota et al., 2004). Clearly, further work is needed to empirically confirm the utility of these descriptions of semantic structure and the mechanisms by which such networks develop over time. However, the recent graph analytic procedures have taken a significant step toward capturing semantic memory within an empirically verified network.

### 2.28.3 Feature Analytic Approaches

An alternative to concepts being embedded within a rich network structure is an approach wherein meaning is represented as a set of primitive features that are used in various combinations to represent different concepts. Of course, this issue (distributed representation of knowledge, by way of features, vs. a localist representation, via a node to concept relationship) is central to attempts to represent and quantify learning and memory in general. We now turn to a review of the feature-based approaches in semantic memory.

The original Collins and Quillian (1969) research generated a great deal of attention, and soon researchers realized that categories reflected more graded structures than was originally assumed. Specifically, some members of categories are good members (ROBIN for BIRD), whereas other members appear to be relatively poor members (VULTURE for BIRD) but are still definitely members of the category (see Battig and Montague, 1969; Rosch, 1973). In addition, there was a clear influence of goodness of an exemplar on response latencies in the sentence verification task described above. Specifically, good exemplars were faster to verify than poor exemplars, referred to as the typicality effect. The Collins and Quillian hierarchical network model did not have any obvious way of accommodating such degrees of category membership.

Smith et al. (1974) took a quite different approach to accommodate the results from the sentence verification task. They rejected the strong assumptions of network theory and proposed a model that emphasized the notion of critical semantic features in representing the meaning of a word. So, for example,

the word BIRD might be represented as animal, two legged, has wings, sings, is small, flies, and so on. There is no hierarchical organization within this model, but concepts reflect lists of critical features. They also distinguished between two classes of features, defining features and characteristic features. Defining features are the necessary features that an exemplar must have to be a member of a category. So, for example, all birds must eat, move, lay eggs, and so forth. On the other hand, characteristic features are features that most, but not all, exemplars have, such as small, flies, sings.

The second important aspect of the Smith et al. (1974) perspective is the emphasis on the decision processes engaged in the classic sentence verification task (see Atkinson and Juola, 1974; Balota and Chumbley, 1984, for similar decision models applied to short-term memory search and lexical decision, respectively). In verifying a sentence such as 'A ROBIN IS A BIRD,' subjects first access all (both defining and characteristic) features associated with ROBIN and all features associated with BIRD. If there is a high degree of overlap in the features, that is, above some criterion, the subject can make a fast 'yes' response. This would be the case in 'A ROBIN IS A BIRD,' since both defining and characteristic features provide a high degree of overlap. On the other hand, some exemplars of a given category may overlap less in characteristic features such as in 'AN OSTRICH IS A BIRD'. Although ostriches are clearly birds, they are not small and do not fly, which are characteristic features of birds. Hence, in such cases, the subject needs to engage in an additional analytic checking process in which only the defining features are compared. This additional check process takes time and so slows response latencies. Hence, the model can naturally capture the typicality effects mentioned above, that is, robins are better exemplars of birds than ostriches, because robins can be verified based on global overlap in features, whereas ostriches must engage the second, more analytic comparison of only the defining features, thereby slowing response latency.

In addition to accounting for typicality effects, the feature analytic model also captured interesting differences in latencies to respond 'no' in the sentence verification task. Specifically, subjects are relatively fast to reject 'A CARP IS A BIRD' compared with 'A BUTTERFLY IS A BIRD.' Carps do not have many overlapping features with birds, and so the subject can quickly reject this item, that is, there is virtually no overlap in features. However, both butterflies and

birds typically have wings, are small, and fly. Hence, the subject must engage the additional check of the defining features for 'BUTTERFLY IS A BIRD,' which ultimately leads to slower response latencies, compared with the sentence 'A CARP IS A BIRD.'

Although there were clear successes of the Smith et al. (1974) feature analytic approach, there were also some problems. For example, the model was criticized for the strong distinction between characteristic and defining features. In fact, McCloskey and Glucksberg (1979) provided a single process random walk model that accommodated many of the same results of the original Smith et al. model without postulating a distinction between characteristic and defining features. According to the random walk framework, individuals sample information across time that supports either a yes or no decision. If the features from the subject and predicate match, then movement toward the yes criterion takes place; if the features do not match, then movement toward the no criterion takes place. This model simply assumed that the likelihood of sampling matching feature information for the subject and predicate is greater for typical members than nontypical members, and therefore the response criterion is reached more quickly for typical than nontypical members, thus producing the influence on response latencies. The distinction between single- and dual-process models is a central issue that pervades much of cognitive science.

A second concern about the Smith et al. (1974) model is that they did not directly measure features but, rather, inferred overlap in features based on multidimensional scaling techniques, in which an independent group of subjects simply rated the similarity of words used in the sentence verification experiment. In this way, one could look at the similarity of the words along an N-dimensional space. Interestingly, Osgood et al. (1957) used a similar procedure to tackle the meaning of words in their classic work on the semantic differential. Osgood et al. found that when subjects rated the similarity across words, and these similarity ratings were submitted to multidimensional scaling procedures, there were three major factors that emerged: Evaluative (good–bad), potency (strong–weak), and activity (active–passive). Although clearly this work is provocative, such similarity ratings do not provide a direct measure of the features available for a concept. So, if there are indeed primitive features, it seems necessary to attempt to more directly quantify these features.

McRae and colleagues have been recently attempting to provide such constraints on feature analytic models (McRae et al., 1997; Cree et al., 1999; Cree and McRae, 2003; McRae, 2004; McRae et al., 2005). The goal here is to develop a feature-based computational model implemented in an attractor network capable of capturing the statistical regularities present in semantic domains. The general notion underlying attractor network models is that knowledge is distributed across units (which might be thought of as features) and that the network settles into a steady pattern of activity that reflects the representation of a concept. The conceptual representations that form the basis of semantic knowledge in the model are derived from feature norming data. To collect norms, groups of participants are asked to list features for a number of concepts (e.g., for DOG, participants might list BARKS, FURRY, CHASES CATS, etc.). McRae and colleagues propose that when participants are asked to list features of various basic-level category exemplars (e.g., DOG and APPLE are basic-level concepts from the superordinate category of MAMMALS and FRUIT, respectively), the resulting lists of features reflect the explicit knowledge people have of these concepts. Importantly, McRae does not claim that the nature of the representation consists of a feature list; rather, he argues, the features are derived from repeated multisensory interactions with exemplars of the concept, and in a feature listing task, subjects temporarily create an abstraction for the purpose of listing features that can be verbally described. Currently, feature norming data are available for 541 concrete objects, representing a wide variety of basic-level concepts. Importantly, the model can account for many empirical observations in semantic tasks, as discussed below.

The major assumption implemented by McRae and colleagues' model is that semantic knowledge, as represented by feature lists, involves the statistical averaging of feature correlations among members of similar categories. Features are correlated if they co-occur in basic-level concepts. For example, HAS FUR is highly correlated with HAS FOUR LEGS, as these two features cooccur in numerous exemplars of the mammal category. However, HAS FUR and HAS WINGS have a low (almost nonexistent) correlation, as these two features do not co-occur frequently. The argument is that individuals are highly sensitive to the regularity of the correlations, which are tapped by semantic tasks. As demonstrated by McRae et al. (1999), the strength of the feature

correlations predicted feature verification latencies in both human subjects and model simulations, with stronger correlations yielding faster response latencies than weaker correlations when the concept name was presented before the feature name (e.g., DOG-FUR). In addition, the correlation strength interacted with stimulus onset asynchrony (the time between the onset of the concept name and the onset of the feature name, SOA). Specifically, the effect of feature correlations was larger at shorter SOAs, with only high correlations predicting response latencies. However, at longer SOAs, even weakly correlated features influenced response times, indicating that, as more time was allowed for the effects of correlated features to emerge, even the more weakly correlated feature-concept pairs benefited from the shared representation. In another series of studies, [McRae et al. \(1997\)](#) reported that the strength of feature correlations predicted priming for exemplars from the living things domain but not for exemplars from nonliving things domains, for which priming was instead predicted by individual features. This finding is consistent with evidence that, compared to living things, nonliving things tend to have a lower degree of correlated features (also see [section 2.28.8.1](#)).

Several interesting extensions of McRae and colleagues' work on the role of features in organizing semantic knowledge have been recently reported. [Pexman et al. \(2002\)](#) examined the role of the number of features associated with a concept and found that items with more features were responded to faster in both naming and lexical decision tasks after a number of other variables known to influence visual word recognition latencies had been factored out. [Pexman et al. \(2003\)](#) reported similar results in a semantic categorization task and in a reading task. These findings were interpreted as supporting the distributed nature of semantic representations in which features are assumed to reflect access to conceptual knowledge, and this information quickly comes on line in isolated word recognition tasks.

In a related vein, there is recent evidence from the categorization literature that categories with richer dimensionalities (i.e., more features and more correlations among features) are easier to learn than categories with fewer dimensions ([Hoffman and Murphy, 2006](#)). Thus, rather than resulting in combinatorial explosions that make learning impossible, rich categories with many features lend themselves well to learning – a finding that is nicely mirrored in how people, even very young children, quickly and reliably learn to recognize and classify objects in the

world. Indeed, it seems that learning to categorize complex objects, which might be quite similar in terms of features, is something most individuals can do reliably and easily. One concern that arises when one examines the richness of the stimuli in the environment, is the potentially infinite number of features that are available to identify a given concept. In fact, critics of feature-based models have argued that the number of possible feature combinations would result in combinatorial explosion, as knowing even a few features of a category could easily result in an enormous number of ways in which the features could be correlated and integrated (see [Murphy, 2002](#), for a discussion). However, as [McRae \(2004\)](#) notes, two points are relevant in addressing this issue. The first is that the feature correlations tend to influence performance largely in implicit tasks – thus reducing the necessity of explaining how an individual can explicitly use the vast amount of information available. The second point is that the feature vectors that underlie semantic representations are generally sparse. In other words, the absence of a specific feature is uninformative, so, for example, knowing that a dog does not have feathers is relatively uninformative. Thus, although feature-based models might not fully capture the richness of the knowledge that individuals have about concepts, they have been useful in advancing research in the field of semantics.

#### 2.28.4 Concept Learning and Categorization

Since semantic memory deals with the nature of representation of meaning, and categories are central to meaning, it is important to at least touch on the area of categorization and how concepts develop. In their classic book, [Bruner et al. \(1956\)](#) emphasized the importance of categorization in organizing what appears to be a limitless database that drives complex human learning and thought. Categorization has been viewed as a fundamental aspect of learning and indeed has been observed early in childhood ([Gelman and Markman, 1986](#)) and in other species such as pigeons ([Herrnstein et al., 1976](#)). Ross ([See Chapter 2.29](#)) provides a much more focused discussion of this topic.

One intriguing question that arises when one considers the content of semantic memory concerns the grain size and structure of the representations. In other words, is there a level at which objects in the

real world are more or less easy to learn and categorize? One possibility is that the world is initially perceived as a continuum in which there are not separate 'things.' Through repeated interactions with verbal labels or other forms of learning, an individual learns how to discriminate separate objects (e.g., Leach, 1964). This approach places the burden on an extensive and demanding learning process. An alternative approach is that the human cognitive system is ideally suited to detect and recognize objects at a specific grain or level. The assumption that the system is biased toward recognizing specific patterns implies that the process of learning the appropriate verbal labels that refer to specific items in the environment is significantly easier. This problem – how very young children learn that when their mother points to a dog and says 'dog' the referent of the phonological pattern in question refers to an entire object, and not to furry things, things of a certain color, or loosely attached dog parts – has been extensively discussed by Quine (1960).

In an elegant series of experiments, Rosch and colleagues (e.g., Rosch et al., 1976) provided empirical evidence that there is indeed a specific level at which categories of objects are represented that contains the most useful amount of information. For example, identifying a given object as a DOG implies that one recognizes that the specific exemplar is a dog, although it may differ from other dogs one has encountered. Simply knowing something is a dog allows one to draw upon a pool of stored knowledge and experiences to infer appropriate behaviors and interactions with the categorized object. However, knowing the object is an animal is not as informative, given the wide variability among animals. For example, interactions with an elephant are likely to be quite different from those one might have with a spider. Conversely, classifying the exemplar as a Collie or as a German Shepherd does not add a significant amount of inferential power for most purposes.

Rosch et al. (1976) argued that at the basic level, categories are highly informative and can be reliably and easily discriminated from other categories. Exemplars of basic-level categories (e.g., DOGS, BIRDS, CARS, etc.) have many attributes in common, tend to be similar in shape and in how one interacts with them, and allow easy extraction of a prototype or summary representation. The prototype can be accessed and serves as a benchmark against which novel exemplars can be compared: Those that are highly similar to the prototype will be quickly and easily classified as members of the category.

Exemplars that differ from the prototype will be recognized as less typical members of a category (e.g., penguins are quite different from many other birds). Hence, typicality effects fall quite nicely from this perspective.

Historically, there has again been some tension between abstract prototype representation and more feature-based approaches. Consider the classic work by Posner and Keele (1968). Although cautious in their interpretation, these researchers reported evidence suggesting that prototypes (in this case a central tendency of dot patterns) were naturally abstracted from stored distortions of that prototype, even though the prototype was never presented for study. They also found that variability across a sufficient number of distortions was critical for abstracting the prototype. These results would appear to support the notion that there is a natural tendency to abstract some representation that is a central tendency of exemplars that share some common elements. So, 'dogness' may be abstracted from the examples of dogs that one encounters. This could suggest that there is indeed a unified representation for dogness.

An alternative approach is that there is nothing unique about these central tendencies but, rather, such representations reflect the similarity of the episodically stored representations in memory. This is a particularly important observation because it suggests that there is a blending of different types of memories, that is, categorical information is simply decontextualized episodic memories. Consider, for example, the classic MINERVA model developed by Hintzman (1986, 1988). In this computational model, each episodic experience lays down a unique trace in memory, which is reflected by a vector of theoretical features. There is no special status of category representations or hierarchical structure. Rather, categorization occurs during retrieval when a probe (the test item) is presented to the system, and the feature vector in the probe stimulus is correlated with all the episodically stored traces. The familiarity of a test probe is a reflection of the strength of the correlations among elements in memory. Because the schema overlaps more with multiple stored representations, that is, it is the central tendency, it will produce a relatively high familiarity signal or strength in a cued recall situation. The importance of the Hintzman approach is that there is no need to directly store central tendencies, as they naturally arise out of the correlation among similar stored traces in the feature vectors. Moreover, as Hintzman argues, there is no need to propose a

qualitative distinction between episodic and semantic memories, because both rely on the same memory system, that is, a vast storehouse of individual feature-based episodic traces.

The notion that categories are a reflection of similarity structure across memory traces and can be generated during retrieval clearly has some appeal. Indeed, Barsalou (1985) demonstrated the importance of ad hoc categories that seem to be easily generated from traces that do not inherently have natural category structure; for example, what do photographs, money, children, and pets have in common? On the surface, these items do not appear to be similar – they do not belong to the same taxonomic category, nor do they share many features. However, when given the category label “things to take out of the house in the case of a fire,” these items seem to fit quite naturally together because our knowledge base can be easily searched for items that are in the house and are important to us. As Medin (1989) has argued, similarity depends on the theoretical frame that a participant uses to guide a search of memory structures. There appears to be an unlimited number of ways in which similarity can be defined, and hence similarity discovered. For example, lichen and squirrels are similar if one is interested in specifying things in a forest. This brings us to the remarkable context dependency of meaning, and the possibility that meaning is not defined by the stimulus *per se* but is a larger unit involving both the stimulus and the surrounding context. The word DOG in the context of thinking about house pets compared with the word DOG in the context of guard dogs or drug-sniffing dogs probably access quite different interpretations, one in which the focus is on companionship, furri-ness, and wagging tails, the other in which the more threatening aspects of dogness, such as sharp teeth, are accessed. One might argue that the context activates the relevant set of features, but even this is difficult until one has sufficient constraint on what those features actually are.

### 2.28.5 Grounding Semantics

In part because of the difficulties in defining the critical features used to represent meaning and potential problems with the tractability of prototypes of meaning, several researchers attempted to take novel approaches to the nature of the representations. There are two general approaches that we review in this section. First, because of the increase in

computational power, there has been an increased reliance on analyses of large-scale databases to extract similarities across the contexts of words used in various situations. This perspective has some similarity to the exemplar-based approach proposed by Hintzman (1986, 1988) and others described earlier. In this sense, meaning is grounded in the context in which words and objects appear. The second approach is to consider the perceptual motor constraints afforded by humans to help ground semantics, that is, the embodied cognition approach. We review each of these in turn.

#### 2.28.5.1 Grounding Semantics in Analyses of Large-Scale Databases

This approach attempts to directly tackle the poverty of the stimulus problem when considering the knowledge that humans have acquired. Indeed, since the days of Plato, philosophers (and more recently psychologists and linguists) have attempted to resolve the paradox of how humans can acquire so much information based on so little input. Specifically, how is it that children learn so much about the referents of words, when to use them, what their syntactic class is, what the relations among referents are, and so on, without explicit instruction? Some have argued (e.g., Pinker, 1994) that the poverty of the stimulus is indeed the reason one needs to build in genetically predisposed language acquisition devices. However, recent approaches to this issue (e.g., Latent Semantic Analysis, or LSA, Landauer and Dumais, 1997; Hyperspace Analogue to Language, or HAL, Burgess and Lund, 1997) have suggested that the stimulus input is not so impoverished as originally assumed. One simply needs more powerful statistical tools to uncover the underlying meaning and the appropriate database.

In an attempt to better understand how rich the stimulus is when embedded in context, Landauer and Dumais (1997) analyzed large corpora of text that included over 4.6 million words taken from an English encyclopedia, a work intended for young students. This encyclopedia included about 30 000 paragraphs reflecting distinct topics. From this, the authors constructed a data matrix that basically included the 60 000 words across the 30 000 paragraphs. Each cell within the matrix reflected the frequency that a given word appeared in a given paragraph. The data matrix was then submitted to a singular value decomposition, which has strong similarities to factor analysis to reduce the data matrix to

a limited set of dimensions. Essentially, singular value decomposition extracts a parsimonious representation of the intercorrelations of variables, but, unlike factor analysis, it can be used with matrices of arbitrary shape in which rows and columns represent the words and the contexts in which the words appear. In this case, the authors reduced the matrix to 300 dimensions. These dimensions reflect the intercorrelations that arise across the words from the different texts. So, in some sense the 300 dimensions of a given word will provide information about the similarity to all other words along these 300 dimensions, that is, the degree to which words co-occur in different contexts. The exciting aspect from this data reduction technique is that by using similarity estimates, the model actually performs quite well in capturing the performance of children acquiring language and adults' performance on tests based on introductory textbooks. In this way, the meaning of a word is being captured by all the past experiences with the word, the contexts with which that word (neighbors) occurs, the contexts that the neighbors occur in, and so on.

The remarkable success of LSA, and other similar approaches such as HAL (Burgess and Lund, 1997), provides a possible answer to the poverty of the stimulus problem, that is, when considering the context, the stimulus is indeed very rich. In the past, we simply have not been able to analyze it appropriately. Moreover, the model nicely captures the apparent contextual specificity of meaning in that meaning is defined by all the contexts that words have appeared in and hence will also be constantly changing ever so slightly across subsequent encounters. Finally, the model is indeed quite important because it does not rely on a strong distinction between semantic and episodic memory since it simply reflects past accumulated exposure to language. In this sense, it has some similarity to the Hintzman (1986) model described above.

### 2.28.5.2 Grounding Semantics in Perceptual Motor Systems

An alternative approach that has been receiving considerable recent attention is that meaning can be grounded in perceptual-motor systems (e.g., Barsalou, 1999). Briefly, this perspective is part of the embodied cognition approach that posits that the cognitive system of any organism is constrained by the body in which it is embedded (Wilson, 2002). Thus, cognition (in this case meaning) is not viewed as being

separable from perceptual, motor, and proprioceptive systems; rather, it is through the interactions of these systems with the environment that cognition emerges. Furthermore, the type of representations that an organism will develop depends on the structure of the organism itself and how it exists in the world. This approach has its roots in Gibson's (1979) ecological psychology, as it is assumed that structures in the environment afford different interactions to different organisms. It is through repeated interactions with the world that concepts and knowledge emerge. Importantly, the very nature of this knowledge retains its connections to the manner in which it was acquired: Rather than assuming that semantic memory consists of amodal, abstract representations, proponents of embodied approaches argue that representations are grounded in the same systems that permitted their acquisition in the first place (Barsalou et al., 2003).

According to the modality-specific approaches to knowledge, a given concept is stored in adjacent memory systems rather than being abstracted. For example, in Barsalou's (1999) account, knowledge is stored in perceptual symbol systems that emerge through repeated experience interacting with an object or an event. Briefly, Barsalou assumed that when a percept is encountered, selective attention focuses on context-relevant aspects of the percept and allows modal representations to be stored in memory. Repeated interactions with similar events or members of the same category result in the formation of a complex, multimodal representation, and a simulator emerges from these common representations. Simulators are the basic unit of the conceptual system and consist of a frame (which is somewhat similar to a schema), the purpose of which is to integrate the perceptual representations. Simulators provide continuity in the system. Importantly, the representations that are stored include not only modal, perceptual information (e.g., sounds, images, physical characteristics) but also emotional responses, introspective states, and proprioceptive information. Retrieving an exemplar or remembering an event is accomplished by engaging in top-down processing and activating the targeted simulator. Importantly, a given simulator can yield multiple simulations, depending on the organism's goal, the context, and the relevant task demands. For example, different simulations for DOG are possible, such that a different pattern of activity will occur if the warm and furry aspect of dog is relevant or whether the aspect of being a guard or police dog is relevant. Of course, this nicely captures the context sensitivity of meaning. Barsalou (1999) argues that perceptual

symbol systems are as powerful and flexible as amodal models, as they are able to implement a complete conceptual system (see also [Glenberg and Robertson, 2000](#)).

Evidence in support of modal approaches to semantics can be found in both behavioral and cognitive neuroscience studies. We briefly review some of this evidence here, although a full review of the neuroscience literature is beyond the scope of this chapter (*See* Chapter 3.07 for further discussion of this area). For example, there is evidence from lesion studies that damage to the pathways supporting a specific modality results in impaired performance in categorization and conceptual tasks that rely on that same modality. Specifically, damage to visual pathways generally results in greater impairment in the domain of living things, which tend to rely heavily on visual processes for recognition. Conversely, damage to motor pathways tends to impair knowledge of artifacts and tools, as the primary mode of interaction with these items is through manipulation (see [Martin, 2005](#)). Consistent with the lesion data, neuroimaging studies indicate that different regions of the cortex become active when people process different categories. Regions adjacent to primary visual areas become active when categories such as animals are processed (even if the presentation of the stimulus itself is not in the visual modality), whereas regions close to motor areas become active when categories such as tools are processed. These findings have been interpreted as consistent with the hypothesis that people run perceptual-motor simulations when processing conceptual information ([Barsalou, 2003](#)).

[Pecher et al. \(2003\)](#) reported evidence from a property verification task indicating that participants were faster in verifying properties in a given modality (e.g., BLENDER-loud) after verifying a different property for a different concept in the same modality (e.g., LEAVES-rustling) than when a modality switch was required (e.g., CRANBERRIES-tart). Pecher et al. argued that the switch cost observed was consistent with the hypothesis that participants ran perceptual simulations to verify the properties (in this case sounds) rather than accessing an amodal semantic system. In a subsequent study, [Pecher et al. \(2004\)](#) observed that when the same concept in a property verification task was paired with two properties from different modalities, errors and latencies increased when verifying the second property. Pecher et al. interpreted this finding as indicating that recent experiences with a concept influence the simulation of the concept. Importantly, researchers have argued

that such results are not simply a result of associative strength (i.e., priming) nor of participants engaging in intentional imagery instructions ([Barsalou, 2003](#); [Solomon and Barsalou, 2004](#)).

Although the results summarized above are compelling and are supportive of the hypothesis that sensory-motor simulations underlie many semantic tasks, the majority of these studies have examined tasks such as property verification and property generation. The question thus arises of whether the results are somehow an artifact of the task demands, and specifically whether these results reflect the structure of the semantic memory system or whether subjects are explicitly retrieving information as they notice the relations embedded within the experimental context. [Glenberg and Kaschak \(2002\)](#) extended the evidence for embodiment effects to a novel series of tasks that do not appear as susceptible to task demand effects. In these experiments, participants read a brief sentence and judged whether the sentence made sense or not. The critical sentences contained statements that implied motion either toward the participant (e.g., “Nancy gave you the book”) or away from the participant (e.g., “You gave the book to Nancy”). Participants responded by moving their hand toward themselves or away from themselves. Glenberg and Kaschak found what they called the action-sentence compatibility effect: When the required response was consistent with the movement implied in the sentence, participants were faster than when the implied motion and the actual physical response were inconsistent. These data appear most consistent with the view that when processing language, people relate the meaning of the linguistic stimulus to action patterns.

### 2.28.6 Measuring Semantic Representations and Processes: Insights from Semantic Priming Studies

As described above, there have been many empirical tools that have been used to provide insights into the nature of semantic memory. For example, as noted earlier, some of the early work by [Osgood et al. \(1957\)](#) attempted to provide leverage on fundamental aspects of meaning via untimed ratings of large sets of words and multidimensional scaling techniques. With the advent of interest in response latencies, researchers turned to sentence verification tasks that dominated much of the early work in the 1970s

and 1980s. Although this work has clearly been influential, the explicit demands of such tasks (e.g., explicitly asking subjects to verify the meaningfulness of subject-predicate relations) led some researchers in search of alternative ways to measure both structure and retrieval processes from semantic memory. There was accumulating interest in automatic processes (LaBerge and Samuels, 1974; Posner and Snyder, 1975) that presumably captured the modular architecture of the human processing system (Fodor, 1983), and there was an emphasis on indirect measures of structure and process. Hence, researchers turned to semantic priming paradigms.

Meyer and Schvaneveldt (1971) are typically regarded as reporting the first semantic priming study. In this study, subjects were asked to make lexical decisions (word-nonword decisions) to pairs of stimuli. The subjects' task was to respond yes only if both strings were words. The interesting finding here was that subjects were faster to respond yes when the words were semantically related (DOCTOR NURSE), compared with when they were unrelated (BREAD NURSE). This pattern was quite intriguing because subjects did not need to access the semantic relation between the two words to make the word/nonword decisions. Hence, this may reflect a relatively pure measure of the underlying structure and retrieval processes, uncontaminated by explicit task demands. Moreover, the development of this paradigm was quite important because researchers thought it may tap the spreading activation processes that was so central to theoretical developments at the time.

The research on semantic priming took a significant leap forward with the dissertation work of Neely (1977), who used a framework developed by Posner and Snyder (1975) to decouple the attentional strategic use of the prime-target relations from a more automatic component. In this study, subjects only made lexical decisions to the target, and subjects were given explicit instructions of how to use the prime information. For example, in one condition, subjects were told that when they received the prime BODY, they should think of building parts (Shift condition), whereas in a different condition, subjects were told that when they received the category prime BIRD, they should think of birds (Nonshift condition). Neely varied the time available to process the prime before the target was presented by using SOAs ranging from 250 to 2000 ms. The important finding here is that the instructions of what to expect had no influence on the priming effects at the short SOA (i.e., priming occurred if the

prime and target had a semantic relationship, independent of expectancies), but they did have a large effect at the long SOA, when subjects had time to engage an attentional mechanism (i.e., the priming effects were totally dependent on what subjects were told to expect, independent of the preexisting relationship). Hence, Neely argued that the short SOA data reflected pure automatic measures of the semantic structure and retrieval processes and could be used as a paradigm to exploit the nature of such semantic representations.

A full review of the rich semantic priming literature is clearly beyond the scope of the present chapter (see Neely, 1991; Lucas, 2000; Hutchison, 2003, for excellent discussions of the methodological and theoretical frameworks). However, it is useful to highlight a few issues that have been particularly relevant to the current discussion. First, there is some controversy regarding the types of prime-target relations that produce priming effects. For example, returning to the initial observation by Meyer and Schvaneveldt (1971), one might ask if DOCTOR and NURSE are related because they share some primitive semantic features or are simply related because they are likely to co-occur in the same contexts in the language. Of course, this distinction reflects back on core assumptions regarding the nature of semantic information, since models like LSA might capture priming between DOCTOR and NURSE, simply because the two words are likely to cooccur in common contexts. Researchers have attempted to address this by selecting items that vary on only one dimension (see, e.g., Fischler, 1977; Lupker, 1984; Thompson-Schill et al., 1998). Here, semantics is most typically defined by category membership (e.g., DOG and CAT are both semantically related and associatively related, whereas MOUSE and CHEESE are only associatively related). Hines et al. (1986), De Mornay Davies (1998), and Thompson-Schill et al. (1998) have all argued that priming is caused by semantic feature overlap because of results indicating priming only for words that shared semantic overlap versus those did not, when associative strength was controlled. However, Hutchison (2003) has recently argued that the studies that have provided evidence for pure semantic effects (i.e., while equating for associative strength), actually have not adequately controlled for associative strength based on the Nelson et al. (1998) free-association norms. Clearly, equating items on one dimension (associative strength or semantic overlap) while manipulating

the other dimension is more difficult than initially assumed. In this light, it is interesting to note that two recent review papers have come to different conclusions regarding the role of semantics in semantic priming based on such item selection studies. Lucas (2000) argued that there was clear evidence of pure semantic effects, as opposed to associative effects, whereas Hutchison (2003) was relatively more skeptical about the conclusions from the available literature.

Balota and Paul (1996) took a different approach to the meaning versus associative influence in priming via a study of multiple primes, instantiating the conditions displayed in Table 1. As one can see, the primes were either both related, first related, second related, or both unrelated to the targets, and the targets could either be homographic words with distinct meanings (e.g., ORGAN) or a nonhomographic words (e.g., STRIPES). As one can see, the primes were related to the targets at both the semantic and associative level for the nonhomographs (e.g., LION and STRIPES are both related to TIGER at the associative and semantic level), but for the homographs the primes were related to the targets at only the associative level (e.g., PIANO and KIDNEY are only related to ORGAN at the associative level, since KIDNEY and PIANO are different meanings of ORGAN). Thus, one could compare priming effects in conditions in which primes converged on the same meaning of the target (nonhomographs) and priming effects where the primes diverged on different

meanings (homographs). The results from four experiments indicated that the primes produced clear additive effects, that is, priming effects from the single related prime conditions nicely summated to predict the priming effects from the double related prime conditions for both homographs and nonhomographs, suggesting that the effects were most likely a result of associative level information. Only when subjects directed attention to the meaning of the word, via speeded semantic decisions, was there any evidence of the predicted difference between the two conditions. Hence, these results seem to be supportive of the notion that standard semantic priming effects are likely to be the result of associative-level connections instead of meaning-based semantic information. Of course, the interesting theoretical question is how much of our semantic knowledge typically used is caused by overlap in the contexts in which items are stored as opposed to abstracted rich semantic representations.

Hutchison (2003) notes two further findings that would appear to be supportive of associative influences underlying semantic priming effects. First, one can find evidence of episodic priming in lexical decision and speeded word naming tasks. In these studies, subjects study unrelated words such as (CITY-GRASS) and are later presented prime-target pairs in a standard lexical decision study. The interesting finding here is that one can obtain priming effects in such studies, compared to an unrelated/unstudied pair of words (see McKoon and Ratcliff, 1979). Thus, the semantic priming effects obtained in word recognition tasks can also be produced via purely associative information that develops within a single study exposure. However, it should be noted here that there is some question regarding the locus of such priming effects and that one needs to be especially cautious in making inferences from the episodic priming paradigm and the role of task-specific strategic operations (see, e.g., Neely and Durgunoglu, 1985; Durgunoglu and Neely, 1987; Spieler and Balota, 1996; Pecher and Raajmakers, 1999; Faust et al., 2001).

The second pattern of results that Hutchison (2003) notes as being critical to the associative account of semantic priming effects is mediated priming. In these situations, the prime (LION) is related to the target (STRIPES) only through a non-presented mediator (TIGER). So, the question is whether one can obtain priming from LION to STRIPES, even though these two words appear to be semantically unrelated. Although de Groot (1983) failed to obtain mediated priming effects in the

**Table 1** Prime-target conditions from the Balota and Paul (1996) multiprime study

<i>Nonhomographs</i>			
<i>Condition</i>	<i>Prime 1</i>	<i>Prime 2</i>	<i>Target</i>
Related-related	LION	TIGER	STRIPES
Unrelated-related	FUEL	TIGER	STRIPES
Related-unrelated	LION	SHUTTER	STRIPES
Unrelated-unrelated	FUEL	SHUTTER	STRIPES
<i>Homographs</i>			
<i>Condition</i>	<i>Prime 1</i>	<i>Prime 2</i>	<i>Target</i>
Related-related	KIDNEY	PIANO	ORGAN
Unrelated-related	WAGON	PIANO	ORGAN
Related-unrelated	KIDNEY	SODA	ORGAN
Unrelated-unrelated	WAGON	SODA	ORGAN

Balota DA and Paul ST (1996) Summation of activation: Evidence from multiple primes that converge and diverge within semantic memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 22: 827-845.

lexical decision task, Balota and Lorch (1986) argued that this may have resulted from the task-specific characteristics of this task. Hence, Balota and Lorch used a speeded pronunciation task and found clear evidence of mediated priming. Evidence for mediated priming has also now been found in versions of the lexical decision task designed to minimize task-specific operations (e.g., McNamara and Altarriba, 1988; McKoon and Ratcliff, 1992; Sayette et al., 1996; Livesay and Burgess, 1998). Of course, it is unclear what semantic features overlap between LION and STRIPES, and so these results would appear to be more consistent with an associative network model, in which there is a relationship between LION and TIGER and between TIGER and STRIPES, along with a spreading activation retrieval mechanism (see McKoon and Ratcliff, 1992; Chwilla and Kolk, 2002, for alternative accounts of the retrieval mechanism).

In sum, although the semantic priming paradigm has been critical in measuring retrieval mechanisms from memory, the argument that these effects reflect amodal semantic representations that are distinct from associative information has some difficulty accommodating the results from multiprime studies, episodic priming studies, and mediated priming studies. As noted earlier, there are available models of semantic memory (e.g., Burgess and Lund, 1997, HAL; Landauer and Dumais, 1997, LSA) and categorization (e.g., Hintzman, 1996, MINERVA) that would strongly support the associative contributions to performance in such tasks and, indeed, question the strong distinction between semantic and episodic memory systems. Hence, this perspective predicts a strong interplay between the systems. We now turn to a brief discussion of the evidence that directly addresses such an interplay.

### 2.28.7 The Interplay Between Semantics and Episodic Memory

Memory researchers have long understood the influence of preexisting meaning on learning and memory performance (see Crowder, 1976, for a review). Indeed, in his original memory manifesto, Ebbinghaus (1885) was quite worried about this influence and so purposefully stripped away meaning from the to-be-learned materials by presenting meaningless trigrams (KOL) for acquisition. Of course, semantics has penetrated episodic memory research in measures of category clustering (see Bousfield, 1953; Cofer et al., 1966;

Bruce and Fagan, 1970), retrieval-induced inhibition (see Anderson et al., 1994), and release from proactive interference (see Wickens, 1973), among many other paradigms. Indeed, the interplay between preexisting knowledge and recall performance was the centerpiece of the classic work by Bartlett (1932). Researchers realized that even consonant–vowel–consonant trigrams were not meaningless (see Hoffman et al., 1987). At this level, one might even question what it would mean to episodically store in memory totally meaningless information.

One place where researchers have attempted to look at the interplay between semantic and episodic structures is within the episodic priming paradigm described earlier. In these studies, participants receive pairs of unrelated words for study and then are later given prime-target pairs that have either been paired together or not during the earlier acquisition phase. For example, Neely and Durgunoglu (1985) investigated the influence of studying previous pairs of words and word–nonword combinations on both lexical decision performance and episodic recognition performance (also see Durgunoglu and Neely, 1987). Although there were clear differences between the tasks in the pattern of priming effects (suggesting dissociable effects across the two systems), there were also some intriguing similarities. For example, there was evidence of inhibition at a short prime-target SOA (150 ms) in both the episodic recognition task and the lexical decision task from semantically related primes that were in the initial studied list but were not paired with the target. It appeared as if this additional semantic association had to be suppressed in order for subjects to make both the episodic recognition decision and the lexical decision. The finding that this effect occurred at the short SOA also suggests that it may have been outside the attentional control of the participant.

The power of preexisting semantic representations on episodic tasks has recently taken a substantial leap forward with the publication of an important paper by Roediger and McDermott (1995), which revisited an earlier paper published by Deese (1959). This has now become known as the DRM (after Deese, Roediger, and McDermott) paradigm. The procedure typically used in such studies involves presenting a list of 10–15 words for study (REST, AWAKE, DREAM, PILLOW, BED, etc.) that are highly related to a critical nonpresented item (SLEEP). The powerful memory illusion here is that subjects are just as likely to recall (or recognize) the critical nonpresented item (SLEEP) as

items that were actually presented. Moreover, when given remember/know judgments (Tulving, 1985), participants often give the critical nonpresented item remember judgments that presumably tapped detailed episodic recollective experience. It is as if the strong preexisting semantic memory structure is so powerful that it overwhelms the episodic study experience.

It should not be surprising that many of the same issues that have played out in the semantic memory research have also played out in the false memory research. Indeed, one model in this area is the activation monitoring (AM) framework (e.g., Roediger et al., 2001a), which suggests that subjects sometimes confuse the activation that is produced by spreading activation that converges on the critical nonpresented item (much akin to the Collins and Loftus, 1975) with the activation resulting from the study event. This framework attempts to keep separate the episodic and semantic systems but also shows how such systems can interact. In contrast, Arndt and Hirshman (1998) have used the Hintzman (1986) MINERVA framework to accommodate the DRM effect by relying on the similarity of the vectors of the individually stored words and the critical nonpresented items. As noted above, the MINERVA framework does not make a strong distinction between episodic and semantic systems. Moreover, the MINERVA model is more a feature-based model, whereas the AM framework *a priori* would appear more akin to a prototype model, but no strong claims have been made along this dimension. A further distinction between the AM framework and the MINERVA approach concerns the relative contributions of backward associative strength (BAS, or the probability that a list item will elicit the target, or critical lure, on a free-association task) and forward associative strength (FAS, or the probability the critical lure will elicit a list item in such a task). According to AM accounts, the critical variable is expected to be BAS, as the activation flows from the list items to the critical lure. However, according to MINERVA, FAS should be more important, as the similarity between the probe (i.e., the critical lure) and the stored episodes (i.e., the list items) should be a more powerful determinant of memory performance. Results from a multiple regression analysis reported by Roediger et al. (2001b) indicated that, in the DRM paradigm, BAS was the better predictor, thus supporting the AM framework. (We thank Roddy Roediger for pointing this out.)

The question of the nature of the representation (i.e., associative vs semantic) underlying these powerful memory illusions has also been studied. For example, Hutchison and Balota (2005) recently utilized the summation paradigm developed by Balota and Paul (1996), described earlier, to examine whether the DRM effect reflects meaning-based semantic information or could also be accommodated by primarily assuming an associative level information. Hence, in this study, subjects studied lists of words that were related to one meaning or related to two different meanings of a critical nonpresented homograph (e.g., the season meaning of FALL or the accident meaning of FALL). In addition, there were standard DRM lists that only included words that were related to the same meaning of a critical nonpresented word (e.g., such as SLEEP). Consistent with the Balota and Paul results, the results from both recall and recognition tests indicated that there was no difference in the pattern of false memory for study lists that converged on the same meaning (standard DRM lists) of the critical nonpresented items and lists that diverged on different meanings (homograph lists) of the critical nonpresented items. However, when subjects were required to explicitly make gist-based responses and directly access the meaning of the list, that is, is this word related to the studied list, there was clear (and expected) difference between homograph and nonhomograph lists. Hutchison and Balota argued that although rich networks develop through strategic use of meaning during encoding and retrieval, the activation processes resulting from the studied information seem to primarily reflect implicit associative information and do not demand rich meaning-based analysis.

There is little doubt that what we store in memory is a reflection of the knowledge base that we already have in memory, which molds the engram. Hence, as noted earlier, semantic memories may be episodic memories that have lost the contextual information across time because of repeated exposures. It is unlikely that a 50-year-old remembers the details of hearing the Rolling Stones' "Satisfaction" for the first time, but it is likely that, soon after that original experience, one would indeed have vivid episodic details, such as where one was, who one was with, and so on. Although this unitary memory system approach clearly has some value (e.g., McKoon et al., 1986), it is also the case that there is some powerful evidence from cognitive neuroscience that supports a stronger distinction.

## 2.28.8 Representation and Distinctions: Evidence from Neuropsychology

Evidence for the distinction of multiple memory systems has come from studies of patients with localized lesions that produce strong dissociations in behavior. For example, the classic case of HM (see [Scoville and Milner, 1957](#)) indicated that damage to the hippocampus resulted in impairment of the storage of new episodic memories, whereas semantic knowledge appeared to be relatively intact (but see [MacKay et al., 1998](#)). Hence, one might be overly concerned about the controversy from the behavioral studies regarding the distinct nature of semantic and episodic memory systems. However, there are additional neuropsychological cases that are indeed quite informative about the actual nature of semantic representations.

### 2.28.8.1 Category-Specific Deficits

There have now been numerous cases of individuals who have a specific lesion to the brain and appear to have localized category-specific deficits. For example, there have been individuals who have difficulty identifying items from natural categories (e.g., animals, birds, fruits, etc.) but have a relatively preserved ability to identify items from artificial categories (e.g., clothing, tools, furniture). At first glance, such results would appear to suggest that certain categories are represented in distinct neural tissue that have or have not been disrupted by the lesion. Such a pattern may also be consistent with a localized representation of meaning instead of a distributed feature-based representation in which all concepts share vectors of the same set of primitive features.

Unfortunately, however, the interpretation of impaired performance on natural categories and intact performance on artificial categories has been controversial. For example, such deficits could occur at various stages in the information flow from discriminating visually similar items (e.g., [Riddoch and Humphreys, 1987](#)) to problems retrieving the appropriate name of an object (e.g., [Hart et al., 1985](#)). Such accounts do not rely on the meaning of the categories but suggest that such deficits may reflect correlated dimensions (e.g., difficulty of the visual discrimination) that differ between natural and artificial categories. In this light, it is particularly important

that there have been cases that have shown the opposite pattern. For example, [Sacchett and Humphreys \(1992\)](#) reported an intriguing case that shows disruption of the performance on artificial categories and body parts but relatively preserved performance on natural categories. They argued that one possible reason for this pattern is that this individual had a deficit in representing functional features, which are more relevant to artifactual representations and body parts than natural categories such as fruits and vegetables. Whatever the ultimate explanation of these category-specific deficits, this work has been informative in providing a better understanding of how members within categories may differ on distinct dimensions.

In a similar vein, one hypothesis that has been suggested to explain domain differences in category-specific deficits is the sensory/functional hypothesis ([Warrington and McCarthy, 1987](#); [Farah and McClelland, 1991](#); [Caramazza and Shelton, 1998](#)). According to this proposal, natural categories such as animals depend heavily on perceptual information (especially on visual discriminations) for identification and discrimination. Conversely, functional information is more important for recognition of artifacts, such as tools. Thus, damage to regions of sensory cortex is expected to result in selective impairment of natural kinds, whereas damage to regions in or adjacent to motor cortex would result in impairment in artifacts. Although compelling, this view is not endorsed by all researchers. [Caramazza and colleagues](#), in particular, have argued that the sensory/functional hypothesis fails to account for some of the patterns of deficits observed and some of the finer-grain distinctions. In particular, it is difficult for this model to account for the selective sparing or impairment of fruits and vegetables, body parts, and musical instruments that have been reported (see [Capitani et al., 2003](#), for a recent review). Thus, the question of whether and how the type of knowledge that is most critical for supporting the representation of a particular domain is involved in category-specific deficits remains open.

To address this controversy, [Cree and McRae \(2003\)](#) extended the sensory/functional hypothesis to include a broader range of types of knowledge. They developed a brain region taxonomy that included nine different forms of knowledge, including sensory/perceptual in all modalities (vision, taste, audition, etc.), functional, and encyclopedic. Encyclopedic features included information about items such as LIVES IN AFRICA for ELEPHANT – in other words,

information that likely was learned and not experienced directly. Cree and McRae then developed a nine-dimensional representation for the 541 concepts for which they had norming data and estimated the salience of each type of knowledge for each object and each category. In a series of cluster analyses, Cree and McRae found that the knowledge types nicely predicted the tripartite distinction between living things, artifacts, and fruits and vegetables reported in several neuropsychological case studies. In addition, Cree and McRae examined several distributional statistics, including the number of distinguishing and distinctive features and similarity to obtain a measure of confusability (i.e., the extent to which a given concept might be confused with another concept from the same category). The categories they examined did appear to be differentially sensitive to these measures, and the implemented model reflected patterns of impairment observed in patients. They concluded that knowledge type does underlie the organization of conceptual representations and that selective impairment in a particular brain region involved in maintaining such knowledge can result in the observed patterns of impairment in patients with category-specific deficits. Although many questions remain, it is clear that evidence from individuals with category-specific deficits has provided considerable insight into both the nature of category representation and the underlying neural representations.

### 2.28.8.2 Semantic Dementia

The most common form of dementing illness is dementia of the Alzheimer type (DAT). However, there is also a relatively rare and distinct dementia, referred to as semantic dementia (SD), which overlaps with DAT in features such as insidious onset and gradual deterioration of comprehension and word-finding ability. SD is a variant of frontal temporal dementia and typically involves one or both of the anterior portions of the temporal lobes. The consensus criteria for SD (Hodges et al., 1992) include impairment in semantic memory causing anomia, deficits in both spoken and written word comprehension, a reading pattern consistent with surface dyslexia (i.e., impairment in reading exception words such as PINT but preserved reading of regular words and nonwords that follow standard spelling to sound rules, such as NUST), impoverished knowledge about objects and/or people with relative sparing of phonological and syntactic components of speech output, and perceptual and nonverbal

problem solving skills. These individuals are often quite fluent, but their speech is relatively limited in conveying meaning. They are particularly poor at picture naming and understanding the relations among objects. For example, the Pyramids and Palm Trees test developed by Howard and Patterson (1992) involves selecting which of two items (e.g., a palm tree or a fir tree) is most similar to a third item (e.g., a pyramid). Individuals with SD are particularly poor at this task and so would appear to have a breakdown in the representations of the knowledge structures.

An interesting dissociation has been made between SD individuals and DAT individuals. In particular, Simons et al. (2002) recently found a double dissociation, wherein individuals with SD produced poorer picture naming than individuals with DAT; however, individuals with SD produced better performance than individuals with DAT on a later episodic recognition test of these very same pictures (also see Gold et al., 2005). Clearly, the selective impairment across these two groups of participants is consistent with distinct types of information driving these tasks. Of course, one must be cautious about the implications even from this study, because it is unlikely that either task is a process-pure measure of episodic and semantic memory (see Jacoby, 1991), but clearly these results are very intriguing.

Recently, Rogers et al. (2004) proposed a model of semantic memory that maintains strong connections to modality-specific systems in terms of both inputs and outputs and has been particularly useful in accommodating the deficits observed in SD. This model has some interesting parallels to Barsalou's (1999) proposal, in that it assumes that semantic memory is grounded in perception and action networks. In addition, like the model proposed by McRae et al. (1997), Rogers et al. suggest that the system is sensitive to statistical regularities, and these regularities are what underlie the development of semantics. The particular contribution of Rogers et al.'s model, however, is that although semantic representations are grounded in perception-action modality-specific systems, the statistical learning mechanism allows the emergence of abstract semantic representations. Importantly, inputs to semantics are mediated by perceptual representations that are modality specific, and as a result, the content of semantic memory relies on the same neural tissue that supports encoding. However, different from Barsalou and colleagues' account, Rogers et al. do

suggest that there is a domain-general, abstracted representation that emerges from cross-modal mappings. Thus, although the system relies on perceptual inputs, the abstract representations can capture cross-modality similarities and structures to give rise to semantic memory.

Rogers et al. (2004) implemented a simple version of their model using a parallel distributed-processing approach in which visual features provided the perceptual input and are allowed to interact in training with verbal descriptors through a mediating semantic level. Importantly, the semantic representations emerge through the course of training as the network learns the mappings between units at the visual and verbal levels. The units the model was trained on consisted of verbal and visual features generated in separate norming sessions. Once training was complete, several simulations were reported in which the model was progressively damaged in a way that was thought to mimic varying levels of impairment observed in individuals with SD. Overall, the model nicely captured the patterns of performance of the patients. Specifically, one pattern often observed in SD is a tendency to overregularize conceptual knowledge. For example, individuals might refer to all exemplars of a category using the superordinate label or a single label that is high in frequency (e.g., calling a DOG an ANIMAL or a ZEBRA a HORSE). This is possibly a result of the progressive failure in retrieving idiosyncratic information that serves to distinguish exemplars, such that only the central tendency (e.g., a prototype or most typical exemplar) remains accessible. Thus, less common items might take on the attributes of higher-frequency exemplars. The model displayed similar patterns of generalization as the SD individuals, a finding explained in terms of changes in attractor dynamics that resulted in the relative sparing of features and attributes shared by many exemplars but a loss of more distinctive features. This model provides an interesting account of semantic memory and the deficits observed in individuals with SD, one in which both perceptually based information and abstracted representations interact to give rise to knowledge of the world.

### 2.28.9 Neuroimaging

Investigations into the nature of semantic memory have benefited from recent advances in technology that allow investigators to examine online processing of information in the human brain. For example,

positron emission technology (PET) and functional magnetic resonance imaging (fMRI) allow one to measure correlates of neural activity *in vivo* as individuals are engaged in semantic tasks (see Logothetis and Wandell, 2004). Although a full review of the substantial contributions of neuroimaging data to the questions pertaining to semantics is beyond the scope of this chapter (See Chapter 3.07 for a review), we briefly examine some of the major findings that have helped constrain recent theorizing about the nature and locus of semantic representations. Two major brain regions have been identified through neuroimaging studies: left prefrontal cortex (LPC) and areas within the temporal lobes, particularly in the left hemisphere.

The first study to report neuroimaging data relevant to semantic memory was conducted by Petersen et al. (1988), who used PET techniques to localize activation patterns specific to semantic tasks. Subjects were asked to generate action verbs upon presentation of a concrete object noun, and activity during this task was compared with the activity occurring during silent reading of the words. Petersen et al. reported significant patterns of activity in LPC, a finding that has since been replicated and extended to other types of attributes. Martin et al. (1995) extended this work to show that the specific attribute to be retrieved yielded different patterns of activation. Specifically, the locus of activation involved in attribute retrieval tends to be in close proximity to the neural regions that are involved in perception of the specific attributes. Thus, retrieval of visual information, such as color, tends to activate regions adjacent to the regions involved in color perception, whereas retrieval of functional information results in activation of areas adjacent to motor cortex. These findings mesh nicely with the perceptual/motor notions of representation in semantic memory reviewed above (e.g., Barsalou, 1999; Rogers et al., 2004). In addition, Roskies et al. (2001) reported that not only were regions in lateral inferior prefrontal cortex (LIPC) preferentially active during tasks that required semantic processing, but specific regions were also sensitive to task difficulty. Thus, it appears that frontal regions are involved both in the active retrieval from semantic memory and in processing specific semantic information.

Many researchers have suggested, however, that although frontal regions are involved in semantic retrieval, the storage of semantic information is primarily in the temporal regions (see Hodges et al., 1992). Indeed, another area that has been implicated in semantic processing is in the ventral region of the

temporal lobes, centered on the fusiform gyrus, and especially in the left hemisphere. This area shows significant activation during word reading and object naming tasks, indicating it is not sensitive to the stimulus form but to the semantic content therein (see [Martin, 2005](#), for a review). Furthermore, within this area, different subregions become more or less activated when subjects view faces, houses, and chairs (e.g., [Chao et al., 1999](#)), suggesting that different domains rely on different regions of neural tissue. This, of course, could be viewed as consistent with the category-specific deficits reviewed above. However, as noted by [Martin and Chao \(2001\)](#), although peak activation levels in response to objects from different domains reflect a certain degree of localization, the predominant finding is a pattern of broadly distributed activation throughout the ventral temporal and occipital regions, which is consistent with the idea that representations are distributed over large cortical regions.

Recently, [Wheatley et al. \(2005\)](#) reported data from a semantic priming study using fMRI that also converges on the notion of perceptual motor representations of meaning. Subjects silently read related, unrelated, or identical word pairs at a 250-ms SOA while being scanned. The related pairs consisted of category members that were not strongly associatively related (e.g., DOG-GOAT, but see the discussion above regarding the difficulty of selecting such items). Given the relatively fast SOA and that no overt response was required, [Wheatley et al.](#) argued that any evidence for priming should be a reflection of automatic processes. Consistent with other evidence that indicates there are reliable neural correlates of behavioral priming that were evidenced by reduced hemodynamic activity ([Wiggs and Martin, 1998](#); [Mummary et al., 1999](#); [Rissman et al., 2003](#); [Maccotta and Buckner, 2004](#)), [Wheatley et al.](#) found decreased activity for identity pairs and a slightly smaller, but still significant, decrease for related pairs relative to the unrelated pairs condition. Importantly, [Wheatley et al.](#) were able to compare patterns of activation as a function of domain. Consistent with proposals by [Barsalou \(1999\)](#), they found that objects from animate objects yielded more activity in regions adjacent to sensory cortex, whereas manipulable artifacts resulted in greater activity in regions adjacent to motor cortex. These findings were taken as evidence that conceptual information about objects is stored, at least in part, in neural regions that are involved in perception and action.

Although the [Wheatley et al. \(2005\)](#) study used a task that was likely to minimize strategic processing, one question that remains to be addressed is whether the automatic and strategic processes involved in semantic priming tasks (see earlier discussion) can also be dissociated in neural tissue. In a recent study, [Gold et al. \(2006\)](#) reported that several of the brain regions previously implicated in processing during semantic tasks are differentially sensitive to the automatic and strategic processes involved in lexical decision tasks. In three experiments, [Gold et al.](#) manipulated prime target relatedness, SOA, and whether primes and targets were orthographically or semantically related. Long and short SOAs were intermixed in scanning runs to assess the relative contributions of strategic and automatic processes (see [Neely, 1991](#)). A comparison of orthographic and semantic priming conditions was included to determine whether any areas were particularly sensitive to the two sources of priming or whether priming effects are more general mechanisms. The results clearly indicated that different regions responded selectively to different conditions. Specifically, midfusiform gyrus was more sensitive to automatic than strategic priming, but only for semantically related primes, as this region did not show reduced activity for orthographic primes. Four regions were more sensitive to strategic than automatic priming, two in left anterior prefrontal cortex and bilateral anterior cingulate. Even more intriguing, the two regions in LIPC were further dissociated: The anterior region showed strategic semantic facilitation, as evidenced by decreased activity, relative to a neutral baseline, whereas the posterior region showed strategic semantic inhibition, or increased activity, relative to the neutral baseline. In addition, the medial temporal gyrus showed decreased activation concurrently with the anterior LIPC, supporting previous claims that these regions show greater activation in tasks that are more demanding of strategic processes but reduced activation when the strategic processes are less demanding ([Wagner et al., 2000](#); [Gold et al., 2005](#)). In sum, it appears that the behavioral dissociations between automatic and strategic processes in priming tasks are also found in the neuroimaging data. The complexity of the patterns of activation involved in semantic tasks appears to indicate that the retrieval and storage of semantic information is indeed a distributed phenomenon that requires the coordination of a wide array of neural tissue.

### 2.28.10 Development and Bilingualism

Although we have attempted to provide a review of the major issues addressed in semantic memory research, there are clearly other important areas that we have not considered in detail because of length limitations. For example, there is a very rich area of developmental research addressing the acquisition of meaning in children (see Bloom, 2000, for a comprehensive review), along with work that attempts to capture the nature of semantic memory in older adulthood (see, e.g., Balota and Duchek, 1989). Of course, we touched upon these issues earlier when discussing how the small world networks of Steyvers and Tenenbaum (2005) develop over time, along with the work by Rosch (1975) on the development of categorization. Given that meaning is extracted from interactions with the environment, the developmental literature is particularly important to understand how additional years of experience mold the semantic system, especially in very early life. There are many interesting connections of this work to topics covered earlier in this chapter. For example, regarding the influence of preexisting structures on false memory, it is noteworthy that young children (5-year-olds) are more likely to produce phonological than semantic false memories, whereas older children (around 11 years and older) are more likely to produce the opposite pattern (see Dewhurst and Robinson, 2004). Possibly, this is a natural consequence of the development of a rich semantic network in early childhood that lags behind a more restricted phonological system.

Another very active area of research involves the nature of semantic representations in bilinguals (see Francis, 1999, 2005, for excellent reviews). For example, researchers have attempted to determine whether there is a common semantic substrate that is amodal, with each language having specific lexical level representations (e.g., phonology, orthography, syntax, etc.) that map onto this system. This contrasts with the view that each language engages distinct semantic level representations. Although there is still some controversy, the experimental results seem more consistent with the assumption that the semantic level is shared across languages, at least for skilled bilinguals. Evidence in support of this claim comes from a diverse range of tasks. For example, in a mixed language list, memory for the language of input is generally worse than memory for the concepts

(e.g., Dalrymple-Alford and Aamiry, 1969). In addition, one finds robust semantic priming effects by translation equivalents (words in different languages with the same meaning, e.g., DOG in English and HUND in German), which is consistent with at least a partially shared semantic representation (e.g., de Groot and Nas, 1991; Gollan et al., 1997).

### 2.28.11 Closing Comments

The nature of how humans develop, represent, and efficiently retrieve information from their vast repository of knowledge has for centuries perplexed investigators of the mind. Although there is clearly considerable work to be done, recent advances in analyses of large-scale databases, new theoretical perspectives from embodied cognition and small world networks, and new technological developments allowing researchers to measure, *in vivo*, brain activity, are making considerable progress toward understanding this fundamental aspect of cognition.

### References

- Albert R and Barabasi AL (2000) Topology of evolving networks: Local events and universality. *Phys. Rev. Lett.* 85: 5234–5237.
- Anderson JR (2000) *Learning and Memory: An Integrated Approach*, 2nd edn. Pittsburgh: Carnegie Mellon.
- Anderson JR and Bower GH (1973) *Human Associative Memory*. Washington DC: Hemisphere Press.
- Anderson MC, Bjork RA, and Bjork EL (1994) Remembering can cause forgetting: Retrieval dynamics in long-term memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 20: 1063–1087.
- Arndt J and Hirshman E (1998) True and false recognition in MINERVA2: Explanations of a global matching perspective. *J. Mem. Lang.* 39: 371–391.
- Atkinson RC and Juola JF (1974) Search and decision processes in recognition memory. In: Krantz DH, Atkinson RC, et al. (eds.) *Contemporary Developments in Mathematical Psychology: I Learning, Memory, and Thinking*, pp. 243–293. San Francisco: Freeman.
- Baddeley A (2000) Short-term and working memory. In: Tulving E and Craik FIM (eds.) *The Oxford Handbook of Memory*, pp. 77–92. Oxford: Oxford University Press.
- Balota DA and Chumbley JI (1984) Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *J. Exp. Psychol. Hum. Percept. Perform.* 10: 340–357.
- Balota DA and Duchek JM (1989) Spreading activation in episodic memory: Further evidence for age-independence. *Q. J. Exp. Psychol.* 41A: 849–876.
- Balota DA and Lorch RF (1986) Depth of automatic spreading activation: Mediated priming effects in pronunciation but not in lexical decision. *J. Exp. Psychol. Learn. Mem. Cogn.* 12: 336–345.

- Balota DA and Paul ST (1996) Summation of activation: Evidence from multiple primes that converge and diverge within semantic memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 22: 827–845.
- Balota DA, Cortese MJ, Sergent-Marshall S, Spieler DH, and Yap MJ (2004) Visual word recognition of single syllable words. *J. Exp. Psychol. Gen.* 133: 336–345.
- Barsalou LW (1985) Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 11: 629–654.
- Barsalou LW (1999) Perceptual symbol systems. *Behav. Brain Sci.* 22: 577–609.
- Barsalou LW (2003) Abstraction in perceptual symbol systems. *Philos. Trans. R. Soc. Lond. Biol. Sci.* 358: 1177–1187.
- Barsalou LW, Simmons WK, Barbey A, and Wilson CD (2003) Grounding conceptual knowledge in modality-specific systems. *Trends Cogn. Sci.* 7: 84–91.
- Bartlett FC (1932) *Remembering*. Cambridge: Cambridge University Press.
- Battig WF and Montague WE (1969) Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. *J. Exp. Psychol. Monogr.* 80: 1–46.
- Bloom P (2000) *How Children Learn the Meanings of Words*. Cambridge MA: MIT Press.
- Bousfield WA (1953) The occurrence of clustering in the recall of randomly arranged associates. *J. Gen. Psychol.* 49: 229–240.
- Bruce D and Fagan RL (1970) More on the recognition and free recall of organized lists. *J. Exp. Psychol.* 85: 153–154.
- Bruner JS, Goodnow JJ, and Austin GA (1956) *A Study of Thinking*. New York: Wiley.
- Burgess C and Lund K (1997) Modeling parsing constraints with high-dimensional context space. *Lang. Cogn. Process.* 12: 177–210.
- Capitani E, Laiacina M, Mahon B, and Caramazza A (2003) What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cogn. Neuropsychol.* 20: 213–261.
- Caramazza A and Shelton JR (1998) Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *J. Cogn. Neurosci.* 10: 1–34.
- Carey S (1978) The child as a word learner. In: Halle M, Bresnan J, and Miller GA (eds.) *Linguistic Theory and Psychological Reality*, pp. 264–293. Cambridge MA: MIT Press.
- Chao LL, Haxby JV, and Martin A (1999) Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nat. Neurosci.* 2: 913–919.
- Chwilla DJ and Kolk HHJ (2002) Three-step priming in lexical decision. *Mem. Cognit.* 30: 217–225.
- Cofer CN, Bruce DR, and Reicher GM (1966) Clustering in free recall as a function of certain methodological variations. *J. Exp. Psychol.* 71: 858–866.
- Collins AM and Loftus EF (1975) A spreading-activation theory of semantic processing. *Psychol. Rev.* 82: 407–428.
- Collins A and Quillian MR (1969) Retrieval time from semantic memory. *J. Verb. Learn. Verb. Behav.* 8: 240–247.
- Cree GS and 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). *J. Exp. Psychol. Gen.* 132: 163–201.
- Cree GS, McRae K, and McNorgan C (1999) An attractor model of lexical conceptual processing: Simulating semantic priming. *Cogn. Sci.* 23: 371–414.
- Crowder RG (1976) *Principles of Learning and Memory*. Hillsdale, NJ: Erlbaum.
- Dalrymple-Alford EC and Aamiry A (1969) Word associations of bilinguals. *Psychon. Sci.* 21: 319–320.
- Deese J (1959) On the prediction of occurrence of particular verbal intrusions in immediate free recall. *J. Exp. Psychol.* 58: 17–22.
- de Groot AMB (1983) The range of automatic spreading activation in word priming. *J. Verb. Learn. Verb. Behav.* 22: 417–436.
- de Groot AMB and Nas GLJ (1991) Lexical representation of cognates and non-cognates in compound bilinguals. *J. Mem. Lang.* 30: 90–123.
- de Mornay Davies P (1998) Automatic semantic priming: The contribution of lexical- and semantic-level processes. *Eur. J. Cogn. Psychol.* 10: 398–412.
- Dewhurst SA and Robinson CA (2004) False memories in children: Evidence for a shift from phonological to semantic associations. *Psychol. Sci.* 15: 782–786.
- Durgunoglu AY and Neely JH (1987) On obtaining episodic priming in a lexical decision task following paired-associate learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 13: 206–222.
- Ebbinghaus H (1885) *On Memory*. Ruger HA and Bussenius CE (trans.) New York: Teacher' College, 1913.
- Faust ME, Balota DA, and Spieler DH (2001) Building episodic connections: Changes in episodic priming with age and dementia. *Neuropsychology* 15: 626–637.
- Farah M and McClelland JL (1991) A computational model of semantic memory impairment: Modality-specificity and emergent category specificity. *J. Exp. Psychol. Gen.* 122: 339–357.
- Fischler I (1977) Semantic facilitation without association in a lexical decision task. *Mem. Cognit.* 5: 335–339.
- Fodor J (1983) *The Modularity of Mind*. Cambridge MA: MIT Press.
- Francis WS (1999) Cognitive integration of language and memory in bilinguals: Semantic representation. *Psychol. Bull.* 125: 193–222.
- Francis WS (2005) Bilingual semantic and conceptual representation. In: Kroll JF and de Groot AMB (eds.) *Handbook of Bilingualism: Psycholinguistic Approaches*, pp. 251–267. Oxford: Oxford University Press.
- Gelman SA and Markman EM (1986) Categories and induction in young children. *Cognition* 23: 183–208.
- Gibson JJ (1979) *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Erlbaum.
- Glenberg AM and Kaschak MP (2002) Grounding language in action. *Psychon. Bull. Rev.* 9: 558–565.
- Glenberg AM and Robertson DA (2000) Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *J. Mem. Lang.* 43: 379–401.
- Gold BT, Balota DA, Cortese MJ, et al. (2005) Differing neuropsychological and neuroanatomical correlates of abnormal reading in early-stage semantic dementia and dementia of the Alzheimer type. *Neuropsychologia* 43: 833–846.
- Gold BT, Balota DA, Jones SJ, Powell DK, Smith CD, and Andersen AH (2006) Dissociation of automatic and strategic lexical-semantics: fMRI evidence for differing roles of multiple frontotemporal regions. *J. Neurosci.* 26: 6523–6532.
- Gollan TH, Forster KI, and Frost R (1997) Translation priming with different scripts: Masked priming with cognates and noncognates in Hebrew-English bilinguals. *J. Exp. Psychol. Learn. Mem. Cogn.* 23: 1122–1139.
- Hart J, Berndt RS, and Caramazza A (1985) Category-specific naming deficit following cerebral infarction. *Nature* 316: 439–440.
- Herrnstein RJ, Loveland DH, and Cable C (1976) Natural concepts in pigeons. *J. Exp. Psychol. Anim. Behav. Process.* 2: 285–302.
- Hines D, Czerwinski M, Sawyer PK, and Dwyer M (1986) Automatic semantic priming: Effect of category exemplar level and word association level. *J. Exp. Psychol. Hum. Percept. Perform.* 12: 370–379.
- Hintzman DL (1986) "Schema abstraction" in a multiple trace memory model. *Psychol. Rev.* 93: 411–428.

- Hintzman DL (1988) Judgements of frequency and recognition memory in a multiple-trace memory model. *Psychol. Rev.* 95: 528–551.
- Hodges JR, Patterson K, Oxbury S, and Funnell E (1992) Semantic dementia: Progressive fluent aphasia with temporal lobe atrophy. *Brain* 115: 1793–1806.
- Hoffman AB and Murphy GL (2006) Category dimensionality and feature knowledge: When more features are learned as easily as fewer. *J. Exp. Psychol. Learn. Mem. Cogn.* 32: 301–315.
- Hoffman RR, Bringmann W, Bamberg M, and Klein R (1987) Some historical observations on Ebbinghaus. In: Gorfein D and Hoffman R (eds.) *Memory and Learning: The Ebbinghaus Centennial Conference*, pp. 57–76. Hillsdale, NJ: Erlbaum.
- Howard D and Patterson K (1992) *The Pyramids and Palm Trees Test*. Bury St. Edmunds: Thames Valley Test Company.
- Hutchison KA (2003) Is semantic priming due to association strength or feature overlap? A micro-analytic review. *Psychon. Bull. Rev.* 10: 785–813.
- Hutchison KA and Balota DA (2005) Decoupling semantic and associative information in false memories: Explorations with semantically ambiguous and unambiguous critical lures. *J. Mem. Lang.* 52: 1–28.
- Jacoby LL (1991) A process dissociation framework: Separating automatic from intentional uses of memory. *J. Mem. Lang.* 30: 513–541.
- Jeong H, Kahng B, Lee S, Kwak CY, Barabasi AL, and Furdyna JK (2000) Monte Carlo simulation of sinusoidally modulated superlattice growth. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 65: 031602–031605.
- LaBerge D and Samuels SJ (1974) Toward a theory of automatic information processing in reading. *Cogn. Psychol.* 6: 293–323.
- Landauer TK and Dumais ST (1997) A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychol. Rev.* 104: 211–240.
- Leach E (1964) Anthropological aspects of language: Animal categories and verbal abuse. In: Lenneberg EH (ed.) *New Directions in the Study of Language*. Cambridge, MA: MIT Press.
- Livesay K and Burgess C (1998) Mediated priming in high-dimensional semantic space: No effect of direct semantic relationships or co-occurrence. *Brain Cogn.* 37: 102–105.
- Logothetis NK and Wandell BA (2004) Interpreting the BOLD signal. *Annu. Rev. Physiol.* 66: 735–769.
- Lucas M (2000) Semantic priming without association: A meta-analytic review. *Psychon. Bull. Rev.* 7: 618–630.
- Lupker SJ (1984) Semantic priming without association: A second look. *J. Verb. Learn. Verb. Behav.* 23: 709–733.
- Maccotta L and Buckner RL (2004) Evidence for neural effects of repetition that directly correlate with behavioral priming. *J. Cogn. Neurosci.* 16: 1625–1632.
- MacKay DG, Burke DM, and Stewart R (1998) HM's language production deficits: Implications for relationships between memory, semantic binding, and the hippocampal system. *J. Mem. Lang.* 38: 28–69.
- Martin A (2005) Functional neuroimaging of semantic memory. In: Cabeza R and Kingstone A (eds.) *Handbook of Functional Neuroimaging of Cognition*, pp. 153–186. Cambridge, MA: MIT Press.
- Martin A and Chao LL (2001) Semantic memory and the brain: Structure and processes. *Curr. Opin. Neurobiol. Cogn. Neurosci.* 11: 194–201.
- Martin A, Haxby JV, Lalonde FM, Wiggs CL, and Ungerleider LG (1995) Discrete cortical regions associated with knowledge of color and knowledge of action. *Science* 270: 102–105.
- McCloskey M and Glucksberg S (1979) Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cogn. Psychol.* 11: 1–37.
- McKoon G and Ratcliff R (1979) Priming in episodic and semantic memory. *J. Verb. Learn. Verb. Behav.* 18: 463–480.
- McKoon G and Ratcliff R (1992) Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *J. Exp. Psychol. Learn. Mem. Cogn.* 18: 1155–1172.
- McKoon G, Ratcliff R, and Dell GS (1986) A critical evaluation of the semantic episodic distinction. *J. Exp. Psychol. Learn. Mem. Cogn.* 12: 295–306.
- McNamara TP and Altarriba J (1988) Depth of spreading activation revisited: Semantic mediated priming occurs in lexical decisions. *J. Mem. Lang.* 27: 545–559.
- McRae K (2004) Semantic memory: Some insights from feature-based connectionist attractor networks. In: Ross BH (ed.) *Psychology of Learning and Motivation*, vol. 45, pp. 41–86. San Diego: Elsevier.
- McRae K, De Sa VR, and Seidenberg MS (1997) On the nature and scope of featural representations of word meaning. *J. Exp. Psychol. Gen.* 126: 99–130.
- McRae K, Cree GS, Westmacott R, and de Sa VR (1999) Further evidence for feature correlations in semantic memory. *Can. J. Exp. Psychol.* (Special issue on models of word recognition) 53: 360–373.
- McRae K, Cree GS, Seidenberg MS, and McNorgan C (2005) Semantic feature production norms for a large set of living and nonliving things. *Behav. Res. Methods Instr. Comput.* 37: 547–559.
- Medin DL (1989) Concepts and conceptual structure. *Am. Psychol.* 44: 1469–1481.
- Meyer DE and Schvaneveldt RW (1971) Facilitation in recognizing words: Evidence of a dependence upon retrieval operations. *J. Exp. Psychol.* 90: 227–234.
- Milgram S (1967) The small world problem. *Psychol. Today* 2: 60–67.
- Miller GA (1990) WordNet: An on-line lexical database. *Int. J. Lexicogr.* 3: 235–312.
- Mummery CJ, Shallice T, and Price CJ (1999) Dual-process model in semantic priming: A functional neuroimaging perspective. *Neuroimage* 9: 516–525.
- Murphy G (2002) *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Neely JH (1977) Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *J. Exp. Psychol. Gen.* 106: 226–254.
- Neely JH (1991) Semantic priming effects in visual word recognition: A selective review of current findings and theories. In: Besner D and Humphreys G (eds.) *Basic Processes in Reading: Visual Word Recognition*, pp. 236–264. Hillsdale, NJ: Erlbaum.
- Neely JH and Durgunoglu AY (1985) Dissociative episodic and semantic priming effects in episodic recognition and lexical decision tasks. *J. Mem. Lang.* 24: 466–489.
- Nelson DL, McEvoy CL, and Schreiber TA (1998) The University of South Florida word association, rhyme, and word fragment norms. <http://www.usf.edu/FreeAssociation/>.
- Osgood CE, Suci GJ, and Tanenbaum PH (1957) *The Measurement of Meaning*. Urbana: University of Illinois Press.
- Pecher D and Raaijmakers JGW (1999) Automatic priming effects for new associations in lexical decision and perceptual identification. *Q. J. Exp. Psychol.* 52A: 593–614.
- Pecher D, Zeelenberg R, and Barsalou LW (2003) Verifying properties from different modalities for concepts produces switching costs. *Psychol. Sci.* 14: 119–124.
- Pecher D, Zeelenberg R, and Barsalou LW (2004) Sensorimotor simulations underlie conceptual representations: Modality-specific effects of prior activation. *Psychon. Bull. Rev.* 11: 164–167.

- Petersen SE, Fox PT, Posner M, Mintun M, and Raichle M (1988) Positron emission topographic studies of the cortical anatomy of single word processing. *Nature* 331: 585–589.
- Pexman PM, Lupker SJ, and Hino Y (2002) The impact of feedback semantics in visual word recognition: Number of features effects in lexical decision and naming tasks. *Psychon. Bull. Rev.* 9: 542–549.
- Pexman PM, Holyk GG, and Monfils M-H (2003) Number of features effects and semantic processing. *Mem. Cognit.* 31: 842–855.
- Pinker S (1994) *The Language Instinct*. New York: Harper Collins.
- Posner MI and Keele SW (1968) On the genesis of abstract ideas. *J. Exp. Psychol.* 77: 353–363.
- Posner MI and Snyder CRR (1975) Attention and cognitive control. In: Solso R (ed.) *Information Processing and Cognition: The Loyola Symposium*, pp. 55–85. Hillsdale, NJ: Erlbaum.
- Quillian M (1968) Semantic memory. In: Minsky M (ed.) *Semantic Information Processing*, pp. 227–270. Cambridge, MA: MIT Press.
- Quine WVO (1960) *Word and Object*. Cambridge, MA: MIT Press.
- Riddoch MJ and Humphreys GW (1987) Visual object processing in optic aphasia: A case of semantic access agnosia. *Cogn. Neuropsychol.* 4: 131–185.
- Rissman J, Eliassen JC, and Blumstein SE (2003) An event-related fMRI investigation of implicit semantic priming. *J. Cogn. Neurosci.* 15: 1160–1175.
- Roediger HL III and McDermott KB (1995) Creating false memories: Remembering words not presented in lists. *J. Exp. Psychol. Learn. Mem. Cogn.* 21: 803–814.
- Roediger HL III, Balota DA, and Watson JM (2001a) Spreading activation and arousal of false memories. In: Roediger HL and Nairne JS (eds) *The Nature of Remembering: Essays in Honor of Robert G. Crowder*, pp. 95–115. Washington DC: American Psychological Association.
- Roediger HL III, Watson JM, McDermott KB, and Gallo DA (2001b) Factors that determine false recall: A multiple regression analysis. *Psychon. Bull. Rev.* 8: 385–407.
- Rogers TT, Lambon RMA, Garrard P, et al. (2004) Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychol. Rev.* 111: 205–235.
- Roget PM (1911) *Roget's Thesaurus of English Words and Phrases*. Available from Project Gutenberg Illinois Benedictine College, Lisle, IL.
- Rosch E (1973) On the internal structure of perceptual and semantic categories. In: Moore TE (ed.) *Cognitive Development and the Acquisition of Language*, pp. 111–144. New York: Academic Press.
- Rosch E (1975) Cognitive representations of semantic categories. *J. Exp. Psychol. Gen.* 104: 192–233.
- Rosch E, Mervis CB, Gray WD, Johnson DM, and Boyes-Braem P (1976) Basic objects in natural categories. *Cogn. Psychol.* 8: 382–440.
- Roskies AL, Fiez JA, Balota DA, and Petersen SE (2001) Task-dependent modulation of regions in left inferior frontal cortex during semantic processing. *J. Cogn. Neurosci.* 13: 1–16.
- Sachett C and Humphreys GW (1992) Calling a squirrel a squirrel but a canoe a wigwam: A category-specific deficit for artefactual objects and body parts. *Cogn. Neuropsychol.* 9: 73–86.
- Sayette MA, Hufford MR, and Thorson GM (1996) Developing a brief measure of semantic priming. *J. Clin. Exp. Neuropsychol.* 8: 678–684.
- Scoville WB and Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Neuropsychiatry* 20: 11–21.
- Simons JS, Graham KS, and Hodges JR (2002) Perceptual and semantic contributions to episodic memory: Evidence from semantic dementia and Alzheimer's disease. *J. Mem. Lang.* 47: 197–213.
- Smith EE, Shoben EJ, and Rips LJ (1974) Structure and process in semantic memory: A featural model for semantic decisions. *Psychol. Rev.* 81: 214–241.
- Solomon KO and Barsalou LW (2004) Perceptual simulation in property verification. *Mem. Cognit.* 32: 244–259.
- Sperling G (1960) The information available in brief visual presentation. *Psychol. Monogr.* 74(11) (Whole No. 498).
- Spieler DH and Balota DA (1996) Characteristics of associative learning in younger and older adults: Evidence from an episodic priming paradigm. *Psychol. Aging* 11: 607–620.
- Steyvers M and Tenenbaum JB (2005) The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cogn. Sci.* 29: 41–78.
- Thompson-Schill SL, Kurtz KJ, and Gabrieli JDE (1998) Effects of semantic and associative relatedness on automatic priming. *J. Mem. Lang.* 38: 440–458.
- Tulving E (1972) Episodic and semantic memory. In: Tulving E and Donaldson W (eds) *Organization of Memory*, pp. 381–403. New York: Academic Press.
- Tulving E (1985) Memory and consciousness. *Can. Psychol.* 26: 1–12.
- Tulving E and Pearlstone Z (1966) Availability versus accessibility of information in memory for words. *J. Verb. Learn. Verb. Behav.* 5: 381–391.
- Wagner AD, Koutstaal W, Maril A, Schacter DL, and Buckner RL (2000) Task-specific repetition priming in left inferior prefrontal cortex. *Cereb. Cortex* 10: 1176–1184.
- Warrington EK and McCarthy R (1987) Categories of knowledge: Further fractionations and an attempted integration. *Brain* 110: 1273–1296.
- Watts DJ and Strogatz SH (1998) Collective dynamics of “small-world” networks. *Nature* 393: 440–442.
- Wheatley T, Weisberg J, Beauchamp MS, and Martin A (2005) Automatic priming of semantically related words reduces activity in the fusiform gyrus. *J. Cogn. Neurosci.* 17: 1871–1885.
- Wickens DD (1973) Some characteristics of word encoding. *Mem. Cogn.* 1: 485–490.
- Wiggs CL and Martin A (1998) Properties and mechanisms of perceptual priming. *Curr. Opin. Neurobiol.* 8: 227–233.
- Wilson M (2002) Six views of embodied cognition. *Psychon. Bull. Rev.* 9: 625–636.