THE CORTICAL DISTRIBUTION AND INTERACTION OF SEMANTIC KNOWLEDGE

by

Robert Frederick Goldberg

B.A., Hofstra University, 1999

M.S., University of Pittsburgh, 2001

Submitted to the Graduate Faculty of
Arts and Sciences at the University of Pittsburgh in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

University of Pittsburgh

2004

UNIVERSITY OF PITTSBURGH

FACULTY OF ARTS AND SCIENCES

This dissertation was presented

by

Robert Frederick Goldberg

It was defended on

October 29, 2004

and approved by

Lawrence Barsalou, Department of Psychology, Emory University

Julie A. Fiez, Departments of Psychology and Neuroscience, University of Pittsburgh

James L. McClelland, Department of Psychology, Carnegie Mellon University

Walter Schneider, Department of Psychology, University of Pittsburgh

Charles A. Perfetti, Departments of Linguistics and Psychology, University of Pittsburgh
Dissertation Director

THE CORTICAL DISTRIBUTION AND INTERACTION OF SEMANTIC KNOWLEDGE

Robert F. Goldberg, PhD

University of Pittsburgh, 2004

Concepts are considered to be the building blocks of human higher-order cognition. Yet theories differ according to how these semantic representations are instantiated within the brain. The amodal characteristics of word meaning imply that this knowledge is stored independent of perceptual experiences. However, mounting evidence suggests that concepts depend upon cortical regions typically ascribed to sensory input. This embodiment of semantic representations through perceptual mechanisms can crucially explain the relationship between the meaning conveyed by words and experience with the associated objects. Across two experiments, this research used functional MRI to examine the role of sensory and prefrontal brain regions while participants verified semantic properties (e.g., *sounds loud*?; *lays eggs*?) of word items. The results show that perceptual properties activate the predicted cortical regions associated with vision, audition, taste and smell, and touch. Increased response times for these perceptual decisions were not associated with increased activity in the identified sensory areas but were associated with increased activity in prefrontal brain regions. In contrast, more abstract semantic decisions led to increased activity in the prefrontal cortex but no such increases were seen for the more difficult decisions. These findings indicate that multiple and widely distributed brain regions used to encode perceptual experiences also support semantic knowledge of those sensory experiences. The prefrontal cortex may represent abstract knowledge and control retrieval with increasing semantic demands for decisions further removed from perceptual experiences.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

In some cases it takes a village, and this dissertation would not have been possible without the wealth of knowledge, mentorship, and support that I received from the Pittsburgh cognitive neuroscience community. I would like to specifically thank: Chuck Perfetti, Walt Schneider, Julie Fiez, Jay McClelland, Jason Chein, Eebie Tricomi, Dave Touretzky, John Opfer, Carl Johnson, Bob Siegler, Carl Olson, Micki Chi, Jonathan Schooler, Erik Reichle, Susan Johnson, Peter Strick, and, from outside Pittsburgh, Larry Barsalou and Sharon Thompson-Schill.

Any attribution of quality and thoroughness to this research is due largely to the invaluable assistance provided by Ikumi Suzuki and Rebecca Smith. The neuroimaging scans benefited from the dedicated assistance provided by Scott Kurdilla at the Brain Imaging Research Center.

This research was supported by the National Science Foundation through the Center for the Neural Basis of Cognition (NSF PGE 9987588) and the Office of Naval Research (ONR 010360) and benefited from a precursor grant from the National Science Foundation (NSF BCS 0121943). I gratefully acknowledge that support.

1. The representation of semantic knowledge: from behavior to brains

The brain is uniquely tuned to extract words from spoken dialogue and events and directly access the underlying meaning. In this regard, Dennett (1984) has called the neural mechanisms a "semantic engine". The converging methodologies of neuropsychology and neuroimaging, constrained by classic and contemporary models of categorization and language processing, suggest how this knowledge may be instantiated within the human brain. By investigating the neural basis of semantic memory, this research examines how the meaning of words is represented by specific brain regions depending on the role of those areas in encoding sensory input and in verbally mediating information about objects.

In focusing on one type of semantic representation, namely concrete objects, and their associated properties (e.g., 'tiger', 'broccoli', 'fur', 'lays eggs', and 'green'), this dissertation research will obviously be limited. But it should be noted that similar conclusions of the functional neuroanatomical architecture could apply across a wide range of other stimuli classes including verbs (Tyler, Russell, Fadili, & Moss, 2001) and numbers (Le Clec'H et al., 2000). For instance, verbs that refer to specific bodily motions are associated with primary motor areas that control movements of the indicated body parts (Hauk et al., 2004). And while this review will be limited to the structure and processing of semantic memory in areas of the cerebral cortex, it should be obvious that subcortical (Middleton & Strick, 1994) and cerebellar (Fiez, 1996) regions are likely to play important, yet still to be clarified, semantic functions.

This introductory chapter will provide a broad overview of how semantic knowledge may be instantiated in the brain through the lens of converging methodologies. Seminal behavioral models indicate how semantic categories may be reduced to their featural elements based on graded relations and the increasing reliance on abstract defining properties when decisions become more difficult. More contemporary theories predict that object knowledge relies on perceptual properties represented in sensory brain regions necessary for encoding experiences with the associated objects. This prediction is supported by studies of neuropsychological patients which indicate how semantic knowledge may be reduced to the constituent features involved based on the roles of dissociable brain regions.

Depending on the specific clustering of perceptual properties within classes of objects, knowledge of some categories, such as living and non-living things, are more likely to be impaired with focal brain damage. For instance, lesions to visual brain areas are more likely to impair knowledge of objects that rely on visual information. Yet the loss of functional information generally spares knowledge of object categories that preferentially depend on visual properties. In this regard, neuropsychological case studies provide evidence for how semantic knowledge can be specifically tied to particular, but widely distributed, brain regions. In contrast, verbally mediated facts about objects appear to depend on prefrontal areas such that even though a patient may not be able to identify an item they can still access abstract information about it.

More recent neuroimaging evidence confirms and extends predictions from the lesion studies. Semantic knowledge can be localized to particular regions of the cortex depending upon the extent to which a specific category or feature represents visual or functional information. Visually-based knowledge activates ventral temporal object recognition areas, whereas facts regarding object functions depend on regions in frontoparietal cortex. Other sensory modalities appear to similarly rely on cortical regions typically ascribed to sensory experiences of objects. The question then arises as to whether perceptual properties of object knowledge in general rely on sensory cortical mechanisms (Allport, 1985) and how more demanding perceptual decisions regarding object knowledge are resolved. In Chapter 2, the first experiment specifically examines these questions by investigating whether a distributed set of sensory brain regions supports perceptual semantic properties associated with the given modalities. In addition, by varying the difficulty of items in each sensory modality, this initial experiment also tested how these areas may be specifically involved as semantic demands increase, and/or are supported by prefrontal brain regions. Studies of the prefrontal cortex indicate the specific role of this region in the controlled retrieval of semantic knowledge and the representation of abstract rules. However, it is not entirely clear whether this prefrontal activity, as suggested from seminal behavioral theories,can be explained through increased semantic abstractness or more general semantic processing demands. In Chapter 3, the second experiment examines the activity of prefrontal regions as the abstractness and demands of semantic decisions were varied. The results of these experiments suggest how semantic knowledge is widely distributed throughout the cerebral cortex and yet highly interactive as the processing demands of semantic decisions increase.

1.1. Behavioral models of semantic knowledge: Reducing categories to features

The componential structure of conceptual knowledge is typically examined by investigating the relationship between object features and knowledge of semantic categories. The classical perspective assumes that categories consist of necessary and sufficient conditions (Katz, 1963). In this regard, the way an item is classified as a member of a category is determined by considering the essential properties of that item and matching those properties to a two-tier feature list of the category. The more recent tradition rejected this explanation in favor of the view that category membership is based on graded property relations (Rips, Shoben & Smith, 1973; Collins & Loftus, 1975; Rosch & Mervis, 1975).

A central assumption of this tradition is that response times and accuracy patterns can be used to determine the organization within and between the categories of object knowledge. Findings of this type led to the view that semantic categories rely on graded relationships based on the overall similarity of features, or lack thereof, among category members. In this regard, this prototype theory of concepts is seen as a rejection of the classical view that categories consist of necessary and sufficient conditions for determining membership. Instead, prototypes are said to reflect the graded similarity relations among category members in which no one feature is crucial for membership, a so-called family resemblance view (Rosch & Mervis, 1975). Prototypes are viewed as a stored representation of the properties that generally tend to represent members of that category. For instance, within the category 'bird', 'robin' and 'cardinal' are considered prototypical members, whereas 'penguin' is not based on the typically shared features of 'has wings', 'can fly', and 'lays eggs'. Properties[1] that are generally weighted as more important for a given category are shared among the most category members. Under response time constraints the more typical members (e.g., 'cardinal' and 'robin') are readily confirmed, as they are said to share the most features with the 'average' member of the category (e.g., birds). In this regard, the similarity between the item and its closest stored prototype is computed based on the similarity of those features. Finally, it is decided whether the new item is 'similar enough' to be called a member of the given category.

There is one clear problem with the prototype approach: How do we decide when an item is 'similar enough' to be a member of a given category, and what happens when an item shares many properties with a category (e.g., Is a 'bat' a 'bird'?) yet cannot be considered a member? That is, how are borderline instances decided? For example, a simple decomposition of the

---

[1] The terms 'feature', 'property', and 'attribute', are treated mostly interchangeable throughout to avoid any possible negative connotation with a particular term. It is not clear whether any meaningful differences can be maintained between these terms with regard to the present discussion.

visual, but also behavioral and scientific, properties often associated with the average ‘bird’ generally predicts the response time findings. Borderline cases (i.e. ‘penguin’ and ‘bat’) take subjects longer to resolve the category membership, in contrast to the more rapid confirmations of the prototypical, or central, members (e.g., ‘robin’) of the category. Yet, penguins do not look like the average ‘bird’ (i.e., they are large, stand upright, and lack feathers) nor do they ‘fly’. And while penguins are adept in and under water, they are still scientifically classified as birds, largely because of their egg-based reproductive cycle that occurs outside of the organism. In contrast, while a bat seems more visually and behaviorally similar to the average ‘bird’, it is correctly classified as a mammal, also largely because of how it reproduces (i.e., gives birth to live young). Prototype models of categorization have trouble with these types of borderline cases because the relevant features must be stored and processed independently of the average category member. The borderline cases would seem to represent exceptions that require defining features to fix category membership, and a prototype representation is insufficient in this regard.

Two-stage models were developed to address this specific difficulty with prototype theories of categorization. Rips, Shoben, & Smith (1973) asked subjects to confirm sentences such as: “A robin is a bird” or “A robin is an animal”, a so called category verification task. They found that subjects were generally faster to confirm the immediate superordinate than the higher superordinate relations, though some interesting exceptions were also apparent. Smith, Shoben, & Rips (1974) posited a two-stage comparison process to account for these differences in verification time. They proposed that the first-stage relies on the overall similarity of features, borrowing from the prototype view, but the second stage would compare only the defining, or relational, features. Interestingly, the second stage could be abbreviated or omitted if there was a very high or low degree of feature similarity. Thus, rapid confirmations (or disconfirmations) could be made on the basis of similarity alone, as in prototype models. The defining features would only be necessary in intermediate, or borderline, cases. In this regard, category membership is fixed through two distinct phases, the first driven by featural similarity, usually based on visual properties (e.g., ‘cardinal’ and ‘hawk’ to the ‘bird’ category), while the second stage relies on defining, or rule-like, features to firmly draw the boundary (e.g., affirming ‘penguin’ while rejecting ‘bat’) around a given category. Many subsequent models have also assumed that two components of processing are required for category verifications though the exact formulations may vary as well the implications of prototypes on theories of concepts (see for example McClosky & Gluckberg, 1979 and Armstrong, Gleitman, & Gleitman, 1983).

A large body of recent literature discusses the distinction between similarity- and rule-based mechanisms but a hybrid model of their interplay is not often the focus of empirical work (Hahn, 1998). Similarity-based processes are seen to account for the reaction times of adults in categorization decisions (Goldstone, 1994; Hampton, 1997). On this view, similarity is based on perceptual mechanisms that compare the particular instance to a summation of a category’s members, perhaps a prototype. However, this view has been criticized on the grounds that children could not acquire a theoretical concept from similarity relations alone (Murphy & Medin, 1985; Keil, 1989), nor is similarity sufficient when category members are exceptions to the general featural tendencies within a given category. Defining features, or even explicit rules, whether based on scientific knowledge or more general relational information, are necessary to account for the wide range of categories apparent within the semantic system. Feature similarity may act as a constraining mechanism in terms of how members of a given category are structured together. That is, the more properties items share in common, the more similar they are, and the greater the likelihood that they belong to the same category and thus have other features in common.

Feature similarity may also be a basic principle of how the brain represents semantic knowledge across sensory modalities. Perceptual properties of object knowledge appear to be represented by the same sensory brain regions associated with the encoding that information. Two items may be considered as sharing similar features if both rely on a common set of sensory brain regions. In this way, category representations may largely arise from perceptual mechanisms responsible representing the information such that perceptually similar categories rely on the same neural substrate. Yet, some categories seem to be well constrained by abstract or defining features (e.g., living things → does it eat and reproduce?) whereas other categories (e.g., furniture) seem to be more a product of the superficial similarity of the constituent, and perceptual, features. In this way, the reduction of category members to particular feature types is a central theme in contemporary theories of semantic memory especially in relating different property types to distinct brain regions.

1.2. Contemporary behavioral models of semantic knowledge representation

More recent theories in the representation of semantic knowledge address how conception and perception may be reunited (Goldstone & Barsalou, 1998). In particular, a perceptually-based proposal for symbol grounding by Barsalou (1999) attempts to model the conceptual system as a product of how the brain encodes and represents sensory information (see also Allport, 1985). By considering sensory experience as a recording system, this perceptual symbol theory accounts for the direct relationship between perceptual features and conceptual knowledge by reducing the former to bottom-up processing in posterior sensory-motor areas of the brain. In later experience, these features are reinitialized by top-down modulatory mechanisms. According to Barasalou (1999), the storage and retrieval of perceptual symbols operates at the level of sensory components that are directly recorded from experience and stored for future use (e.g., 'purr', 'green', and 'hot'). As symbols become organized, with additional weightings from attentional mechanisms, they enable simulations of these components independent of the perceptual experiences that gave rise to them and eventually reference a particular modal representation (i.e., a word or picture). Over time, these representations take on increasing specificity and complexity with symbolic referents gradually standing in for the actual objects so that detailed decisions can be made from memory (e.g., does a bird have: four legs? an antenna?). Category relations emerge under this view by the type of interrelations that occur among the perceptual components, or properties.

The most relevant aspect of this proposal for the current focus is the productivity of symbol representation that it predicts. Since perceptual features of object knowledge are stored in analogous regions of cortex to where those properties were encoded by sensory mechanisms, there would seem to be a limitless base from which all knowledge representation could emerge, without the need for a separate encoding and storage mechanism. In general support of this linking between perception and conception, Glenberg and Robertson (2000) have suggested that words are indexed by affordances to perceptual objects and experiences. In addition, Zwaan and colleagues (Zwaan, Stanfield, & Yaxley, 2002) have found that when subjects are read sentences regarding the location of an animal or an object, the shape of the object changes as a function of its location (e.g., an eagle in the sky or in the nest). Therefore the current behavioral evidence suggests that conceptual knowledge is grounded by perceptual properties, perhaps based on the roles of sensory brain regions in representing this semantic information.

In both the neuropsychological and neuroimaging literatures examining semantic knowledge there has been a consistent blurring between perceptual mechanisms and conceptual knowledge. The ways that properties of objects are instantiated within the cortex demonstrate how category knowledge can arise from overlapping patterns of activity in adjacent brain regions. That is, the similarity between members of the same superordinate category (e.g., animals) is likely based on common set of cortical regions being active, as predicted by Allport (1985). In contrast, the more difficult semantic decisions involve distinguishing between highly similar instances (e.g., tiger and leopard) and uniting subcategories that share very little perceptually in common (e.g., living thing = plants and animals). This breadth and depth of semantic knowledge is likely based upon different types of properties being localized to different regions of cortex. Neuropsychological patients with specific deficits of semantic knowledge provide preliminary evidence for how these types of properties may be organized to particular brain regions.

1.3. The selective impairment of semantic categories through specific properties

When damage is localized to fairly discrete regions of cortex and the loss or preservation of semantic knowledge is observed, there is a strong tendency to associate the impacted area with a role in conceptual representation. Typically such patients demonstrate an inability to name or retrieve information about objects (e.g., tiger) and often via multiple modalities. These patterns of loss or preserved knowledge seem to reflect, on some level, category boundaries, specifically that the representation of an item requires posterior sensory regions of cortex. When these areas are damaged, patients show selective deficits to categories (e.g., animal) that tend to rely upon the particular feature type (e.g., color) represented by that area. However, even when the perceptually-based representation of an item is impaired (e.g., naming it from a description: “What is that smart swimming mammal with fins?”), the patient may still be possible to access abstract information about (e.g., the ‘animal’ was used to find underwater mines during World War II) the item. This distinction is likely due to a dissociable neuroanatomical basis for each type of knowledge. Therefore, case studies provide compelling evidence of the distinct types of properties represented and processed by sensory (i.e., perceptual properties) and prefrontal (i.e., relational and abstract properties) cortical regions.

If a patient with a circumscribed lesion demonstrates impaired performance on one class of categories (e.g., living things) but not another (e.g., non-living things), an existence proof is provided that on some level the underlying neuroanatomy reflects categories boundaries. However this evidence merely provides a starting point in that such cases are relatively rare and are seldom completely overlapping from one patient to the next. The direct evidence from such studies is that semantic knowledge can be explained in terms of the cortical areas necessary for encoding the perceptual properties most relevant for some categories but not others. In this regard, the neuroanatomy is likely to represent much more basic aspects of items in addition to explicit category labels. What exactly is impaired in patients with semantic deficits is the subject of much current debate.

Neuropsychological findings of category distinctions within semantic memory support the relationship between sensory brain mechanisms and conceptual knowledge, as predicted in perceptual symbol systems (Barsalou, 1999). These case studies suggest a neuroantomical distinction between the role of temporal cortical areas in the representation, or content, of semantic knowledge and that of prefrontal areas in processing relational information and abstract rules. Two stage models of category verification seem to require this distinction in the underlying neuroanatomy, especially in order to explain the difference in response times for easy and hard verification judgments within a particular category. Patients with lesions can thus provide a preliminary foundation for understanding the relationship between behavioral models of category structure and processing and the underlying neuroanatomy that gives rise to them.

The neuropsychological literature is filled with case studies that demonstrate a selective loss of semantic memory (for recent reviews, see Saffran & Schwartz, 1994; and Caramazza, 1998) even though such cases are relatively rare (Coltheart, 2001). A distinction between living and non-living things, even if this formulation is inexact, has been the most frequently reported among patients with semantic memory impairments (Mummery, Patterson, Hodges, & Price, 1998). Damage localized to inferomedial temporal cortex, usually resulting from encephalitis of the herpes simplex virus, has been associated with impaired performance for living things (e.g., animals, fruits, and vegetables) relative to a sparing of artifact (e.g., tools, household objects, vehicles) knowledge (Gainotti & Silveri, 1996; Hart & Gordon, 1992). The reverse pattern of category-specific deficits, i.e., non-living things impaired relative to a preservation of living thing knowledge, usually from damage to frontoparietal areas, has also been found (Warrington & McCarthy, 1987; Hills & Caramazza, 1991; Sacchett & Humphreys, 1992). This double dissociation has led to the inference that semantic knowledge is organized on some level by category-specific information.

In the face of category-specific deficits, and more diffuse semantic disturbances when stimulus factors are rigorously controlled, a number of proposals have been advanced to explain the living thing/nonliving dissociation found among patients. The most widely accepted view assumes that there are substantial differences in the saliency of semantic features for particular categories (Mummery et al., 1998) and that the neural instantiation of these features is biased to higher-order cortical areas mostly responsible for encoding them (e.g., visual features in extrastriate regions). The assumption is that living things are distinguished from each other primarily through their visual or perceptual properties (e.g., consider the difference between a tiger and a lion) while artifacts are determined more by functional features (e.g., consider when a jar can be considered a vase). The assumption is that damage to the semantic network differentially affects the neural regions associated with processing visual or functional semantic features and thus the categories that rely on these features (Warrington & Shallice, 1984; Saffran & Schwartz, 1994). This proposal has been termed the perceptual/ functional hypothesis.

Warrington & Shallice (1984) describe four cases, all as a result of herpes simplex encephalic damage to both temporal lobes, in which the patients exhibit much poorer performance in comprehending or producing the names for living things but intact knowledge for non-living things. For example, patient JBR identified the pictures of just 3 of 48 living things yet he was correct on 43 out of 48 inanimate objects. And whereas he defined ‘briefcase’ as a “small case used by students to carry papers”, ‘snail’ was described as “an insect animal”. This category-specific deficit for living things and food items was found to differing degrees in the four patients across a wide range tasks, including verbal descriptions, naming, mimed responses, and picture/word matching. In contrast, Warrington and McCarthy (1983) demonstrate that patient VER shows a relative preservation in the comprehension of animals, plants, and food items (e.g. ‘soup’) but a gross impairment for assorted household objects. Based on these observations, Warrington and McCarthy (1983) propose that the patterns of category-selective sparing and impairment reflects the distinction between perceptual (read: visual) properties more necessary to distinguish cabbage from lettuce and the functional (read: usage) qualities required to differentiate a table from a desk.

Warrington & Shallice (1984) argue that the selective impairment and preservation of semantic categories reflects a distinction between artifacts, that require functional qualities to differentiate similar members while visual properties remain fairly constant, and living things, perhaps natural kinds, which be distinguished by subtle perceptual elements. While many of JBR's difficulties coincided with living items (e.g., flowers and animals), a few impaired categories are clearly non-living even as he was very good with most objects and even body parts. Along these lines, JBR was impaired on musical instruments and precious stones, categories in which items would seem to be differentiated by visual features (consider: 'trumpet' and 'trombone' and 'diamond' and 'ruby') like shape, sounds, and colors. These more anomalous aspects of JBR's performance can be addressed from the perspective of patient YOT (Warrington & McCarthy, 1987) who is shown to exhibit fine-grained categorical deficits and selective sparing in the face of widespread impairments in the capacities for the comprehension and production of prepositional speech. The semantic profile of YOT is almost the mirror reverse of JBR in that her impairments are mostly with small manipulable objects and furniture, but interestingly also body parts, while she demonstrates a selective sparing of animals, flowers, foods, vegetables, along with different types of fabric. Her case provides further evidence for a double dissociation between the knowledge for items differentiated primarily through perceptual features and those distinguished according to their functional attributes.

A number of other studies (e.g., Warrington & McCarthy, 1983; Sacchett & Humphreys, 1992; Moss & Tyler, 2000) have also shown that categories grounded by functional properties can be damaged almost independently of perceptually-based categories. Sacchett and Humphreys

(1992) present patient CW who is selectively impaired at naming and recognizing common artifacts and body parts, even as his performance with animals, fruits, and vegetables, was mostly normal. Knowledge of object use has been shown to be selectively impaired in patients with semantic dementia (Hodges, Bozeat, Lambon Ralph, Patterson, & Spatt, 2000). Furthermore, Buxbaum & Saffran (2002) have found that manipulability, or the knowledge for the type of action used with a given artifact (e.g., hammer - nail), is selectively impaired in apraxics with fronto-parietal lesions. In the behavioral literature, it has been recently proposed that recognizing artifacts require some understanding of a designer's intentions (Bloom, 1996; for discussion see Malt & Johnson, 1998; and Bloom, 1998) or how it should be used. It seems then that the brain honors the extent to which an artifact is reducible to its functional attributes, or its intended abstract use, rather than its shape or color which can vary widely. For example, consider the many variants of a radio from a pocket version to car and home stereos.

Recent computational models implement the perceptual/functional hypothesis that particular conceptual domains differ in the degree and pattern of property correlations (Farah & McClelland, 1991; McRae, de Sa, & Seidenberg, 1997; Devlin, Gonnerman, Andersen, & Seidenberg, 1998; McClelland & Rogers, 2003). For example, using feature ratings from subjects, Farah & McClelland (1991) modeled the general living/non-living distinction as a function of the properties ascribed to each category. When the nodes representing visual-semantic elements were damaged, the model was impaired in producing the names for living things, with impairment increasing with the amount of damage. In contrast, when the functional-semantic units were damaged, the model was unable to correctly name the non-living things. In this regard, the category distinctions are seen to arise from the differential weighting of the constituent properties. More recently, Rogers and McClelland (2004) update and extend this work to show how the mechanisms of the same computational model based on the features of items can account not only for these deficits but also for behavioral response patterns of normal children and adults.

A number of investigators propose alternative hypotheses for these category-specific effects. Rather than as a result of differing degrees of damage to specifically functional or perceptual properties per se, both Humphreys & Forde (2001; Forde, Francis, Riddoch, Rumiati, & Humphreys, 1997) and Moss & Tyler (Moss, Tyler, Durrant Peatfield, & Bunn, 1998; Moss, Tyler, & Jennings, 1997) have both proposed some version of a ‘density’ hypothesis for the structure of semantic categories and their selective breakdown. According to this view, members of a category may share from just a few (e.g., consider: office products) to very many (e.g., consider: four-legged animals) overlapping attributes. These density theories predict that the more properties shared within a given category, the harder it will be to distinguish any one exemplar with damage to that feature set.

For instance, Moss & Tyler (1997) point out that the biological functions that apply to most living things are strongly correlated with common, and largely interacting, perceptual elements like outline shape (Lloyd-Jones & Luckhurst, 2002). Yet the distinctive perceptual features that distinguish one animal from another (e.g., stripes or spots) have no meaningful connection with these functional markers. Keil (1989) has shown that young children have extensive difficulty recognizing that surface features for living things play no role other than aiding in identification, such that if they are shown that the stripes are removed from a tiger, the child will claim it has become a lion. Similarly, Carey (1985) has demonstrated that children must learn to ignore the differences in the outward behavior between plants and animals, to understand the intrinsic biological mechanisms that unite them under the common category. After localized damage, the common perceptual features of living things that support these abstract biological properties remain preserved (Forde et al., 1997) in the face of losses to individual perceptual elements necessary for the identification of individual items (e.g., tiger from leopard). Accordingly, Moss et al. (1998) report that patient RC is nevertheless able to sort living things based on their shared properties even though he demonstrates tremendous difficulty naming the same items.

Artifacts, on the other hand, seem to depend crucially on such subtle, and often times unique, form-function contingencies. That is, the teeth of a saw are just as important for how it looks as for what it does, as are the tines of a fork. In recognizing artifacts, these tighter form-function contingencies of distinctive visual properties to functional attributes leads to better preservation of this knowledge when damage to one type of feature occurs. So even when temporal regions that support shape and color are damaged, the name of the artifact can still be derived from intact functional and relational features.

The cortical segregation of living and non-living things is generally inferred from category-specific effects for animals and fruits on the one hand, and a diverse set of artifacts on the other. However, these abstract, and theory-based (see for instance Carey, 1985), categories are multifaceted, and even vague, and as such, there is very little evidence that either category can be selectively spared or impaired as a whole (Caramazza & Shelton, 1998). More often the observed deficit or sparing is much more selective. In many cases patients demonstrate almost complete recovery from cerebral damage except for a specific category deficit, such as just for animals (Hart & Gordon, 1992; Caramazza & Shelton, 1998), fruits and vegetables (Hart, Berndt, & Caramazza, 1985), or plants (Pietrini et al., 1988). Similarly a selective sparing, in the face of widespread deficits, has been shown for body parts (Shelton et al., 1998). For instance, Hart, Berndt, & Caramazza (1985) report on patient MD who demonstrates a selective impairment for naming fruits and vegetables and yet showed normal performance for animals, trees, assorted household objects, body parts, and clothing stimuli. In addition, MD demonstrates no impairment to other food items in contrast to the general impairment of this broader category found by Warrington and Shallice (1984). The results of this case study are further supported by Hillis & Caramazza (1991) who found a double dissociation between two patients, JJ and PS, where one was more impaired in naming animal items and relatively preserved on other objects, while the reverse was true for the other patient, even a year after their respective cortical insults. Furthermore, Shelton, Fouch, and Caramazza (1998) report that their patient IOC produces relatively spared responses to body parts in contrast to a diffuse impairment in naming and comprehending a wide range of semantic categories, including anomia for animals, fruits and vegetables, clothing, and household objects. These results point to the more exact specificity with which semantic knowledge may be deficient or preserved.

Based on this much finer grain of selective impairments and preserved categories for living and non-living things, Caramazza and Shelton (Caramazza & Shelton, 1998; Caramazza, 1998; Shelton et al., 1998; Shelton & Caramazza, 2001) have proposed that the organization of semantic memory is broadly based on domain-specific, likely innately specified, mechanisms. In support of this domain-based view, they argue that if a category is selectively impaired or spared, a reductionistic feature-based account (i.e., the perceptual/functional or density theories) requires that other categories relying on similar properties, should show complementary performance. However, cases of very specific category effects, as cited above, seem to suggest an alternative explanation. Since the patterns of finer selectivity often favors categories (e.g., animals, plants, fruits and vegetables, and food items) that are inherently important for the survival of an individual, they suggest that evolutionary pressures influenced specific adaptations necessary for recognizing these items. In this regard, category-specific effects would be expected to reflect an underlying organization based on domains of knowledge, rather than on constituent feature sets. For example, it is crucial under survival conditions to recognize animals as either a food source or as a predator, and plant life as a source of medicine or nutrition. According to Caramazza and Shelton (1998), these evolutionary adaptations provide skeletal neural structures, in the form of dedicated circuitry or processing units, around the most salient distinctions in the conceptual world are organized. Under this view, the differences between categories emerge from the distinction between animate and inanimate objects that eventually leads to neuroanatomical differences between the domains of living and nonliving things. In support of this theory, this distinction appears to have its basis in early infancy (for example see Bertenthal, 1993) and so conceptual organization would seem to be grounded innately if it arises directly from it.

This evolutionary-based account has difficulty with the category-specific evidence in a number of ways. First, even if categories are organized at the neural level by domain-based distinctions, the internal structure of a given domain and how it is shaped by experience is still underspecified, as Caramazza (1998) readily admits. This point is made much clearer when we consider the finer-grained cases that this account seeks to explain. Why ‘body parts’ should be preserved (Shelton et al., 1998) independently of ‘animals’ and ‘fruits and vegetables’, and ‘tools’ and ‘kitchen utensils’ for that matter, remains just as underspecified as it was in the perceptual/functional account. Second, more recent studies have been designed to specifically test the domain-specific theory. Borgo & Shallice (2001) report their patient MU to be much more impaired than four other patient controls in visually recognizing not only living things, but also materials (e.g., wood and glass), liquids (e.g., olive oil and alcohol), and substances (e.g., Nutella and ground pepper). The domain-specific theory does not account for these findings as there are no a priori reasons for assuming man-made liquids, materials and substances should be included with the domain of biological kinds, and instead these results would seem to support some form of the perceptual/functional feature-based processing distinction. In order for MU to recognize the sensory-based distinction between, for instance, ‘glass’ and ‘wood’ he had to rely on perceptual attributes that were likely lost as a result of damage to his temporal lobes. In addition, Marques (2002) has shown that in normal subjects, perceptual and functional feature processing is apparent, even when comparisons are made within a domain.

While the developmental evidence does seem to support an early distinction in children between animate and inanimate things, this distinction is more likely based on highly salient movement characteristics rather than on a rudimentary domain-based biological understanding. Instead young children encounter great difficulty when they learn that movement is a mere consequence of life not a causal determinant (Carey, 1985; Keil, 1989). Even adults appear to represent ‘plants’ and ‘animals’ separately, and not under a common domain, based on reaction time evidence in a simple living/non-living category verification task (Goldberg, under review). Thus, there would seem to be little converging support for the domain-specific theory as it is currently formulated.

In general, alternative accounts to the perceptual/functional theory do not capture the category-specific effects as well, nor do they provide a basis of further extension. The perceptual/functional hypothesis remains the most viable, in terms of offering a principled account of why particular categories should be impaired together and why damage to particular regions of cortex are associated with certain deficits. That is, when the temporal lobes are damaged, patients have difficulty with categories that require perceptual, mostly visual, features to distinguish among similar instances (e.g., ‘tiger’ from ‘leopard’). In contrast, if a patient suffers from a lesion to the frontal or parietal lobes, any semantic memory deficit they encounter is more likely to affect categories that rely on functional properties or relational information. The perceptual/functional hypothesis can therefore explain why damage to certain regions of cortex should result in a particular impairment or preservation of semantic knowledge. This theory also offers a preliminary account for how semantic information may be instantiated by the sensory areas responsible for initially encoding that information, as predicted by the perceptual symbol theory, and further blurring the distinction between perception and conception.

Extensions to perceptual/ functional proposal are represented by slightly different variants of density theories with a particular emphasis on whether properties a part of the sensory knowledge of a given item or rely on indirectly learned knowledge about the item. In this regard the neuropsychological literature appears to implicitly draw on a distinction proposed independently by Rips (1995) between representations *of* an object, based on direct sensory experiences, and information *about* the object, based on indirectly learned facts. According to density theories, the representation of a category is damaged or spared based on the form-function feature contingencies that differentiate particular classes of objects (Humphreys & Forde, 2001). Similar members of a given category (e.g., animals) interfere with each other when distinguishing, but superficial, features (e.g., ‘spots’ and ‘stripes’) are disrupted, yet leave intact abstract, and more relational, biological features (Moss et al., 1997) that can still be associated with a mostly preserved category prototype for the gross set of common attributes (e.g., ‘four legs’). That is, the information about a category can remain mostly preserved so long as a basic prototype representation of the category remains. This split between the representation of a category and the information about that category can be explained in terms of the underlying neuroanatomy. For instance, a patient who shows an impaired ability to recognize the perceptual features of items yet who can verbally cite functional and encyclopedic properties of the same items argues for these two types of semantic knowledge subsiding in distinct anatomical regions. In this way, semantic knowledge would seem to be represented not just in terms of the accumulation of sensory and functional features averaged over exemplars, but also through the many relations between and within categories that are coded and simulated through language. The distinction between the representation of and the information about a given object appears to rely on differing functional anatomical substrates, such that one may be damaged independently of the other for a specific object category.

McCarthy & Warrington (1988) report that their patient TOB performed normally when asked to name and describe pictures of living and nonliving things. Yet he was clearly impaired when asked to describe the same living things if the items were presented verbally. For instance to the word ‘dolphin’ he responded, “A fish or a bird.” But when he was presented with a picture of a ‘dolphin’ he responded, “Dolphin lives in water…they are trained to jump up and come out…In America during the war years they started to get this particular animal to go through to look into ships”. Clearly, the label served no real meaning, while performance was spared for artifacts (e.g., lighthouse → “Round the coast, built up, tall building, lights revolve to warn ships). In contrast to this deficit for the verbal recognition of living things, Warrington and McCarthy (1994) describe how patient DRS exhibited a selective deficit for naming pictures of artifacts, even though he could identify them from verbal descriptions and while no deficit was found for either vehicles or animals. Not only are living things and artifacts dissociable according the perceptual features that are involved, but they are also differentiated in terms of the knowledge used to access verbal information about the items. De Renzi and Lucchelli (1994) found their patient Felicia to be impaired on perceptual properties for living items (i.e., animals, fruits, vegetables, flowers, and also foods), but she showed a spared ability to access functional-encyclopedic knowledge of the same items (e.g., for vegetables → used in minestrone?). This selective sparing of abstract relational information is significant because her lesions were most restricted to temporal regions leaving intact the candidate frontal regions necessary for supporting the controlled retrieval of abstract knowledge about the category.

While these diverse results would seem to emphasize the interaction of input modalities, (i.e., sensory vs. verbal) with particular categories, this conclusion is still being debated. Lambon Ralph and colleagues (Lambon Ralph, Graham, Patterson, & Hodges, 1999) studied nine semantic dementia patients when they defined items from either the presentation of a picture or its spoken name. They note that while some patients show superiority for picture-based definitions, this advantage is likely due to the additional amount of information available in the stimuli. Since Lambon Ralph et al. (1999) found no differences between living things and artifacts, they argue that living things cannot be more reliant on perceptual features. And while this finding may seem to refute the perceptual/functional dichotomy discussed earlier, the definition tasks they used require more information about the categories, more so than physical attributes. Therefore it is not surprising that eight out of the nine patients produced more relational (e.g., ostriches live in Africa) than physical (e.g., an ostrich has a long neck) properties. Considering all of the patients suffered from lesions to the inferior aspect of either one or both temporal lobes, the impairment of this physical knowledge was expected. In this regard, while these patients showed no category-specific effects, the abstract encyclopedic knowledge that they rely on to provide verbal definitions draws upon information about these categories presumably based on spared cortical regions in the frontal cortex. This evidence provides further support for the view that semantic knowledge may be dissociated according to abstract relational elements, or information *about* items, and constituent sensory attributes, or representations *of* objects.

The neuropsychological evidence provides two preliminary conclusions regarding what aspects of semantic knowledge are localizable to particular regions of cortex. First, to the extent that categories rely on perceptual or functional properties they may be selectively impaired or spared depending on the location of the damage. Patients with damage to the temporal lobes, usually recovering from encephalitis of the herpes simplex virus, tend to suffer from deficits to their knowledge of natural kinds and associated perceptual, especially visual, properties. As shown later in neuroimaging studies, this region of cortex has also been implicated in object color knowledge as well as more general object recognition processes. In contrast, frontoparietal lesions are more often associated with impaired performance on artifacts and their associated functional properties including how an object is used and what it does.

Of the available theories, the perceptual/ functional hypothesis best describes this distinction and, more generally, how semantic features are localized to regions of cortex necessary for encoding the relevant sensory information, as predicted by perceptual symbol systems. The ventral temporal lobes support the subtle visual features necessary to distinguish among very similar items (e.g., leopard and tiger) most often associated with living things. Artifacts, on the other hand, rely on tighter form-function contingencies (e.g., tines and fork) such that they remain relatively spared with deficits to visual knowledge with the functional attributes aiding in identification. Such intercorrelation of feature types is suggested by density theories that supplement the perceptual/functional hypothesis. The anatomical localization of these two feature types argues for a distribution and representation of semantic knowledge based on the constituent properties of a given category and the sensory-motor areas of cortex necessary for initially encoding that information.

The second preliminary conclusion from the neuropsychological literature is that not only can perceptual properties be selectively damaged or spared, but so too can more abstract relational features. Visual discrimination is required for the fine distinctions between animal exemplars. These form-based differences appear to rely on visual object processing regions in ventrolateral temporal cortices. In contrast, more abstract relational information, coded by verbal cues, seems to be processed, and perhaps represented, by frontal cortical regions. In this regard, the temporal lobes appear to be more directly involved in the representation *of* a category, that is those features that can be directly encoded by perceptual mechanisms, whereas the frontal lobes appear to support semantic processing regarding the features involved in representing information about a given category, or those features necessarily coded by the language system. This dissociation is exemplified by findings that even when patients cannot identify a given item, for example a dolphin, they can still access abstract knowledge regarding aspects such as habitat and behavior. These prefrontal regions involved in such abstract semantic knowledge may rely on specific processing mechanisms or more rule-based representations.

Both of these preliminary findings have been more recently examined in the context of neuroimaging findings of the organization and processing of semantic memory. By comparing the semantic deficits found in neuropsychological patients with neuroimaging activations for mostly the same categories in normal individuals, object concepts are shown to be instantiated in the predicted cortical regions, including the wide distribution of categories depending upon their predominant property type. In this regard, neuroimaging methods more directly localize visual properties to broad networks within inferior temporal regions used in visual object recognition.

In contrast, the prefrontal cortex appears to subserve at least two distinct semantic functions, one involving general semantic processing demands and the other involved in representing abstract and relational semantic information.

1.4. The neuroimaging of semantic knowledge: Localizing content not categories

Neuroimaging methods, such as Positron Emission Tomography (PET) and functional Magnetic Resonance Imaging (fMRI), provide fine-grained spatial (i.e., at the millimeter level) and temporal (i.e., on the order of seconds) resolution in vivo of cognitive operations in the normal brain. These methods enable a direct mapping of the myriad of ways that the brain uses internal codes to process and represent information. Using these methods it is possible to compare the locations of such cognitive operations with evidence from patterns of association and dissociation in the neuropsychological literature (Posner, Petersen, Fox, & Raichle, 1988). With recent advances in understanding the relationship between the hemodynamic signal of local oxygen dependencies and the underlying neural population (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), and aided by cognitive theories, it is now possible to instantiate the complexities of semantic knowledge in terms of how the cortex represents and processes our understanding of the world.

Some of the earliest work in neuroimaging (Petersen, Fox, Posner, Mintun, & Raichle, 1988; Petersen, Fox, Posner, Mintun, & Raichle, 1989) was directed at uncovering how lexical items and the generation of semantic associates rely on different brain regions. Participants were presented with single words representing concrete nouns (e.g., ‘hammer’) and were asked to generate an appropriate verb (e.g., ‘pound’) for each item. Petersen et al. (1989) found that in comparison to when subjects simply read the words, this verb generation task activated a disparate network of brain regions including the prefrontal cortex. This initial evidence has provided the thrust for current investigations of conceptual representations by using neuroimaging methods to more precisely determine the locations responsible for the structure and processing of semantic knowledge.

In a recent review (Cabeza & Nyberg, 2000) of 275 functional neuroimaging studies, 46 studies were shown to address semantic memory retrieval. Most of these studies explore the knowledge of words and objects, and Cabeza and Nyberg differentiate these into two broad groups: those that use either categorization or generation tasks. The categorization tasks ask participants to classify words or pictures into different categories, (e.g., living or non-living?), similar to seminal category verification tasks. Generation tasks, like the verb generation task just discussed, assess the regions involved in word production to a given cue (e.g., stem completion → ti---?). Cabeza and Nyberg also differentiate a smaller set of studies that that use mostly object-naming and word-reading tasks to more directly explore the distinctions between domains (e.g., living and nonliving things) and categories (e.g., animals and tools). Across all of these tasks, semantic memory retrieval is generally associated with greater activation of prefrontal, parietal, and temporal cortices, exactly what we should expect from case studies of category-specific effects. That is, the neuropsychological evidence has shown that damage to these regions produces selective deficits in the retrieval of some categories of items. In addition, these areas map onto traditional language areas (e.g., Broca's area = left inferior frontal gyrus) as well as provide new insights into how semantic knowledge is both organized and processed (Demonet et al., 1992), based on both category-type and the requisite properties involved.

For instance, when subjects are simply asked to monitor a stream of aurally presented animal names (e.g., 'turtle') for whether the animal was both 'native to the United States' and 'used by humans', a widespread, left lateralized network of regions, in contrast to a tone detection task, was uncovered that extends beyond traditional language areas and is consistent across subjects (Binder et al., 1997). The most activated regions included left hemisphere temporoparietal areas adjacent to, but spreading out from, Wernicke's area, including middle and inferior temporal, fusiform, and angular gyri, extensive left prefrontal regions near to Broca's area, as well areas in which the exact semantic functions are unknown, such as the anterior and posterior cingulate gyri, left subcortical regions including thalamus, and the right cerebellum. Somewhat surprisingly, some of these same regions are activated by a standard rest control task that has often been used as a baseline contrast for the cognitive task of interest in neuroimaging studies (Binder et al., 1999) suggesting that the semantic retrieval network is a default obligatory process[2], even during nonlinguistic tasks (Price, Wise, & Frackowiak, 1996).

This broad network of regions must also interact substantially to produce the intuition of unified and yet flexible semantic representations such as category knowledge. For instance, activation patterns found within, for example, temporal cortex directly suggest how the contents of semantic memory may intertwine. Middle temporal gyrus (BA 21) is activated by words and object pictures (Martin, Wiggs, Ungerleider, & Haxby, 1996; Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995; Vandenberghe, Price, Wise, Josephs, & Frackowiak, 1996) suggesting a role in the convergence of semantic information regardless of modality. Even blind and normal subjects are shown to activate the same temporal lobe location when word stimuli are either tactile or visual, relative to nonword letter strings (Buchel, Price, & Friston, 1998). In addition, the left anterior temporal pole has been cited as an important location for the composition of

---

[2] This evidence argues against the use of a simple rest control task in neuroimaging studies, especially as a contrast for linguistic tasks. It seems that subjects tend to engage everyday brain regions when asked to simply 'rest', thus in a direct contrast with the control task, the task of interest may fail to show many relevant regions necessary for performing the cognitive task.

sentence meaning from the interaction of grammar and semantics (Vandenberghe, Nobre, & Price, 2002). Thus a widespread network of cortical regions would appear to be involved in the representation and processing of semantic knowledge.

Based on the neuropsychological evidence, recent neuroimaging investigations of semantic memory have focused on the distinct regions of activity found for different categories and their constituent properties. The general impairment of living things with damage to ventral and inferior temporal regions has been recently clarified with neuroimaging studies that show how visual information may be specifically coded by specialized sub-regions and by more diffuse patterns of cortical activity. In addition, research on the anatomical regions responsible for semantic processing, such as controlled retrieval and selection from competing sources, has highlighted the role of, specifically, the left prefrontal cortex in the representation of abstract, verbally-mediated information.

1.5. Localizing semantic properties to sensory-motor brain regions

The distinction between semantic structures and semantic processing is based on seminal behavioral findings, and this distinction, perhaps because it is so helpful in conceptualizing categories, has been adopted, with a strong push from the neuropsychological findings, by neuroimaging investigations into semantic memory. Notions of semantic structure, in particular, are most amenable to a reductionist perspective on the role of specific types of properties in category membership, and the locations where they can be localized in cortex, mostly the sensory-motor regions necessary for encoding the perception of that information.

The neuropsychological evidence suggests that the differential weighting of property types, (e.g., for animals, subtle visual properties) and their intercorrelations, locates category knowledge into distinct cortical regions (e.g., ventral temporal areas). As the most common dissociation in the neuropsychological literature is between living and nonliving things, this distinction has also been more recently investigated using neuroimaging. In a blocked design participants were asked in separate experiments to sort items, both pictures (Sergent, Ohta, MacDonald, & Pulvermueller, 1992) and words (Price, Moore, Humphreys, & Wise, 1997), into living and nonliving classes. The activity for this gross semantic task was compared to a more general linguistic control, in these studies a letter-sound coherency judgment and syllable counting task, respectively. The results implicated left ventral temporal cortex in this general semantic processing suggesting that participants use visual properties to simply classify items as living or non-living things. However, this result says little about how the different content between the two categories is necessarily represented and accessed.

In a direct comparison between natural kinds and artifacts, Mummery and colleagues (Mummery, Patterson, Hodges, & Wise, 1996) used PET to investigate item retrieval in response to natural kind category labels (e.g., ‘fruit’ and ‘land animals’), manipulable object category labels (e.g., ‘toys’ and ‘weapons’), and initial letters (e.g., ‘S’ and ‘P’). Participants were instructed to articulate as many appropriate exemplars as possible to the different cues during the 20-second scans. They found that this semantic category retrieval task, in contrast to initial letter fluency, activated inferolateral and anteromedial aspects of the left temporal lobe. The direct contrast between semantic categories showed greater activation bilaterally in the anteromedial temporal lobe for the natural kind categories whereas the manipulable objects elicited posterior inferolateral temporal cortex, generally supporting the dissociation of these broad classes that has been found in the lesion literature. This restriction of activity to the temporal lobes appears to indicate that when participants generate category exemplars they rely on visual properties.

These results have been questioned by Devlin and colleagues (Devlin et al., 2002b) in a study first performed using PET and then replicated with fMRI. They initially found differences between living things and artifacts in a lexical decision experiment but these effects disappeared once a more stringent significance level was adopted that corrected for multiple comparisons. They then tested whether this null result was due to the relative ease with which a lexical decision experiment is performed in contrast to a more demanding semantic task. Subjects were asked to read three cue words presented in a sequence and then were to judge whether a fourth word belonged to the same category, either natural kinds or artifacts. This task was then compared with a letter categorization control that asked subjects to view three consecutive trials of differing numbers of the same lowercase letter where the test trial was either the same or a different letter in uppercase font. In using both PET and fMRI, no effect of category was found at the corrected threshold though effects of tools and natural kinds were found at an uncorrected threshold which were somewhat consistent with the previous findings for these categories. The authors argue that while category information is widely distributed in the cortex, the semantic system is undifferentiated by category at the neural level. However, this conclusion may be premature especially considering recent work from some of the same contributors.

Devlin and a mostly separate set of colleagues (Devlin et al., 2002a) examined the inconsistencies of these reports by combining the data from seven prior neuroimaging studies used to investigate category differences into a single multifactorial analysis which crossed category (living vs. artifact) with a range of tasks. Their hypothesis centered on the view that category-specific effects in the neuroimaging literature may be relatively small and task dependent. In support of this prediction, they found reliable category effects in only word retrieval tasks and semantic decisions. In addition, they provided converging support for the findings of Mummery and colleagues (Mummery et al., 1996) as living things were found to activate bilateral anteromedial aspects of the temporal lobes while tools activated a left middle posterior region of temporal cortex. These findings support the view that the content of categories is anatomically dissociable but may be mediated by task conditions. This conclusion is further bolstered by studies that compare more discrete categories that generally share many properties (e.g., animals), the level at which category segregation is likely to occur, as predicted by the perceptual/functional hypothesis. Not surprisingly, the living/nonliving distinction while providing the initial thrust for neuroimaging studies of semantic memory is as underspecified, and perhaps as unhelpful, as it was in the neuropsychological studies. That is, an apple and a tiger would seem to have few features in common though both likely use visual properties for differentiation from near ‘family’ members within their respective superordinate categories.

When using more constrained categories, and thus a greater likelihood for the same sets of properties, neuroimaging methods have been more successful at uncovering the separable anatomical bases for the content of semantic knowledge and the relationship to neuropsychological case studies (Damasio, Grabowski, Tranel, Hichwa, & Damasio, 1996; for discussion see Caramazza, 1996). Replicable activation differences between animals and artifacts have been found (Perani et al., 1995; Martin et al., 1996) with left occipitotemporal regions in the ventral visual stream more active for animals and left prefrontal and premotor regions more active for tools, supporting the specific locations predicted by the neuropsychological case studies. More recently, Chao and Martin (2000) found that simply viewing and naming pictures of tools activated left ventral premotor cortex (BA 6) suggesting a motor analogue for the semantic features of these objects. In addition, recent work (Le Clec'H et al., 2000) has shown that regions adjacent to these areas are active when subjects consider the location of named body parts, further implicating this general region in the representation of categories that require knowledge of movements associated with objects. In general, sensory brain motor areas that are necessary for encoding relevant perceptual information are also associated with the representation of the categories that rely on this information.

More specifically, across a range of tasks, Chao, Haxby, & Martin (1999) show that viewing, matching, and naming animals significantly activates not only occipital visual regions but also elicits bilateral activation in the lateral aspect of the fusiform gyrus including the occipitotemporal sulcus. Importantly, these regions are a part of the upper levels in the visual object processing hierarchy (Haxby et al., 2001). Moore and Price (1999) extend the ventral temporal findings toward the left anterior pole when fruit pictures are named and matched with their verbal labels, similar to the region found by Mummery and colleagues (1996) bilaterally for the much broader natural kind category that included fruit items but with strictly verbal stimuli.

In further support of the localization of semantic properties of objects, Martin, Chao, and colleagues (1995; Chao & Martin, 1999) have found that fusiform areas close to color perception were activated by the generation of color words. These areas are adjacent, and slightly anterior, to regions activated by animals across a wide-range of tasks, suggesting that the fine-grain of perceptual differentiation required to discriminate, for instance a tiger from a leopard, relies on the same cortical regions involved in perceiving this information. In contrast, the generation of action words activated a lateral occipitotemporal region close to motion perception areas. These areas border regions typically activated by tool stimuli suggesting a basis for the manipulation knowledge associated with some artifacts. Even categories that have generally been thought not to rely on visual information show increased activation in these regions when the visual details of these items are probed. Thompson-Shill et al. (1999) have shown that the left fusiform gyrus was more active when participants had to answer visual questions about artifacts (e.g., Are both the front and back end of a submarine approximately the same width?) than when they had to answer non-visual questions (Does a toaster use more electricity than a radio?). In contrast, no differences were found between visual and non-visual questions about living things suggesting that comparable visual semantic features are necessarily retrieved about living things regardless of the task context.

Mummery and colleagues (Mummery et al., 1998) have directly investigated the role of both color and location attributes using similarity judgments across the categories of living and nonliving things. They expected to find ventral temporal regions more active for the color judgments and enhanced dorsal temporal-parietal activation for judging similarity of items based on their typical locations, in line with the now classic delineation of extrastriate processing streams (Ungerleider & Mishkin, 1992). Using a word triad task, subjects were asked to make a decision for which of two items was most similar to a target based on color (e.g., 'flamingo' → 'salmon' or 'cherry') or location (e.g., 'paper'→ 'ruler' or 'rake'). Mummery et al. (1998) found that, as predicted, the left anteromedial temporal cortex showed increased activation for color judgment replicating the role of this region in the representation of visual semantic attributes. In contrast, a dorsal region in the left temporo-occipito-parietal (TOP) junction was more active for judgments involving object locations implicating this area in the representation of semantic properties associated with where an object is typically found. Finally, artifacts showed an enhanced effect in the left posterior middle temporal gyrus, near to areas involved in object motion, replicating studies that associate this region with the representation of action knowledge.

While the vast majority of neuroimaging studies into the content of semantic memory have relied almost exclusively on visual differences between categories, more recent work has looked at how other modalities may rely on sensory areas to process and represent semantic information. Noppeney and colleagues (2002), using an aural presentation of names for semantic features, found anterior and ventral regions of the left temporal lobe active in the processing of sensory-based knowledge including both visual (e.g., ‘blue’) and auditory (e.g., ‘noise’) items. Furthermore when subjects are asked to judge the sound that an object makes, an area in posterior superior temporal cortex is activated that is, not surprisingly, adjacent to auditory association cortex (Kellenbach, Brett, & Patterson, 2001).

To test how cortical regions may be specifically and selectively involved in the representation of semantic categories, we recently examined (Goldberg, Perfetti, & Schneider, 2002) disparate brain regions as participants considered inter-item similarity relations. They were asked to covertly generate ”the most similar” item they could think of to a target item (e.g., ‘duck’ → [‘chicken’]), and then to choose from two items the one that was most similar to the item they had generated (e.g., ‘goose’ or ‘pigeon’). Unbeknownst to the subjects, all trials were drawn from four categories: fruits, birds, body parts, and clothing. These categories were chosen to reflect the neuropsychological dissociation between categories that tend to rely on visual (i.e., fruits and birds) or functional (i.e., clothing and body parts) properties. This semantic task, in contrast to a visuo-motor control condition, activated a left-lateralized network of brain regions including fusiform, inferior and ventral temporal, and anterior cingulate regions. The visual categories, as expected, activated ventral temporal cortex whereas the functional categories tended to drive lateral temporal regions, replicating previous results indicating exactly this dissociation between these types of categories. More interestingly, category-specific regions were found to be active depending upon the distribution of the presumed semantic properties to modalities representative of each set of items and their similarity relations. For example, fruit items selectively activated bilateral regions in orbitofrontal cortex. This region has been shown to represent taste and olfactory information (Rolls, 2001, 2002), exactly what we should expect given the similarity relations among fruit items, in contrast to the other categories tested. In contrast, body parts were shown to specifically activate a region in lateral occipitotemporal cortex near to where pictures of body parts are encoded (Beauchamp et al., 2002). These results generally support the prediction that semantic information is widely distributed to different cortical areas, specifically to sensory-motor regions associated with representing the constituent properties of the given categories. However, it is unclear whether the categories or their constituent properties gave rise to these effects.

Semantic categories and their associated attributes would seem to be localized to discrete neuroanatomical regions depending on the conceptual content. In accord with the perceptual/functional hypothesis, visual categories and their associated properties seem to be represented by inferior and ventral temporal regions that are more generally involved in visual object recognition. In contrast, artifacts and their associated properties have most often been localized to posterior temporal and pre-motor areas, near to regions involved in object motion and manipulation. In this regard, the representation of semantic categories would seem to be reduced to the constituent features of the particular category being represented and the analogous sensory-motor regions of cortex necessary for perceptually encoding this information. However, it is not yet clear if this direct relationship between perceptual experiences of object and semantic knowledge of those items applies across all possible sensory modalities. This predicted relationship between the sensory encoding of object properties and knowledge of them is the subject of the first experiment presented in Chapter 2.

1.6. General and specific contributions of the prefrontal cortex to semantic processing

The prefrontal cortex is most likely responsible for helping us sort through conflicting evidence in semantic memory to determine with conviction whether a peripheral, or borderline, item is a member of the given category. This brain region has been implicated in the laying down of new memories in coordination with temporal cortices, controlled retrieval selection processes in semantic memory, and perceptual classification, as well more general cognitive control in initiating and maintaining goal states. Interestingly, single-unit recording studies have shown how perceptual classification and the learning of rules may be implemented at the neuronal level in prefrontal cortices of monkeys (Miller, Freedman, & Wallis, 2002). Indeed, the prefrontal cortex would seem to be the ideal place to investigate the functional architecture for abstract, and relational, information *about* items in the world as it occupies a far greater portion in the human brain than in other animals (Fuster, 1995). The role of prefrontal cortex, usually left anterior and inferior areas if not a finer segregation (for discussion see Fiez, 1997), in semantic memory retrieval is presently unclear. This region may serve both to represent abstract properties of objects while using these representations in the service of more general semantic processing mechanisms.

The prefrontal cortex has been widely implicated in semantic memory encoding. For instance, left inferior prefrontal cortex shows increased activation during semantic encoding relative to nonsemantic encoding, and the activity was decreased with additional semantic encoding of the same words, but not when the words were reprocessed nonsemantically (Demb et al., 1995). Similar semantic priming effects in inferior prefrontal areas have been widely reported (Bookheimer, 2002) including similar regions for pictures and words (Wagner, Desmond, Demb, Glover, & Gabrieli, 1997) and a task-specific fractionization of this region into anterior and posterior components (Wagner, Koustaal, Maril, Schacter, & Buckner, 2000). In addition, whether previous experiences will be remembered would seem to be predictable from the activation of prefrontal cortex during encoding with activity lateralized based on the modality of the items, with the right hemisphere more active for pictures (Brewer, Zhao, Desmond, Glover, & Gabrieli, 1998) and the left hemisphere more active for word stimuli (Wagner et al., 1998). Based on these findings it should be clear that prefrontal cortex plays a crucial role in learning and memory especially regarding semantic knowledge.

As reported earlier, some of the first neuroimaging investigations of cognition (Petersen et al., 1988, 1989) found left inferior prefrontal activations for when subjects were asked to generate a verb to a target noun in contrast to when nouns were simply repeated. In general, prefrontal activity found in semantic retrieval has usually been lateralized to the left hemisphere, even when the task materials were nonverbal, including objects (Martin et al., 1996; Vandenberghe et al., 1996) and logographic characters (Tan et al., 2001). More recent work shows that subregions of the prefrontal cortex may be segregated based on the modality of the stimulus materials (McDermott, Buckner, Petersen, Kelley, & Sanders, 1999) and by the presented task (Roskies, Fiez, Balota, Raichle, & Petersen, 2001). More specifically, activations generally found in ventrolateral prefrontal cortex during both classification and generation tasks suggest that these areas are related to generic semantic retrieval processes (Cabeza & Nyberg, 2000). In contrast, ventromedial prefrontal cortex is activated by classification tasks suggesting some sort of decision-related function during semantic retrieval and perhaps a role in representing abstract rules necessary for borderline semantic classifications. Lateral activations closer to Broca's area in posterior prefrontal cortex are also found during semantic retrieval processes that could reflect covert or overt articulation during word generation. In addition, as noted earlier, more dorsal regions of prefrontal cortex may be also be involved with the selective processing of tools relative to other objects. Therefore, prefrontal cortex in general may subserve different aspects of semantic retrieval processes including the representation of specific types of knowledge and mechanisms used to guide semantic decisions.

While many studies have focused on the left inferior prefrontal cortex (LIPFC) for its apparent functions in semantic knowledge, its exact role remains unclear (Wagner, Pare-Blagoev, Clark, & Poldrack, 2001). Initially, the LIPFC was thought to subserve semantic retrieval or semantic working memory (Petersen et al., 1989, 1988; Kapur et al., 1994; Demb et al., 1995; Fiez, 1997). In this regard, the LIPFC was hypothesized to be involved in the active 'working with' an item, here specific for semantic information, as seemed to be true of other regions in prefrontal cortex (Kapur et al., 1994) for other stimulus factors in maintaining memories over brief intervals. More recently, it has been proposed that rather than mediate semantic retrieval per se, the LIPFC is specifically involved in the selection of task-relevant responses from among competing alternatives (Thompson Schill, D'Esposito, Aguirre, & Farah, 1997). Thompson-Shill and colleagues (1997) varied selection demands by manipulating the amount of semantic distractors in a given task. For instance, in a verb generation paradigm, subjects covertly produced an associated verb with a target noun. The nouns were varied so that they had many associated verbs ('high-selection': rope → hang, tie, lasso, etc.) or one dominant verb ('low selection': scissors → cut). They found that, in comparison to when the subjects simple read the verbs, activity was greater for both generate conditions in the LIPFC. Critically, when the selection conditions were compared directly, activity in the LIPFC was greater for the high-selection condition. They argue that this result indicates the LIPFC is more generally involved with selection of responses in general from among competing alternatives, rather than simply as a semantic retrieval mechanism.

Wagner and colleagues (Wagner et al., 2001) have recently tested this proposal by examining activity in the LIPFC while selection demands were held constant. Participants were presented with a cue (e.g., candle) and had to decide which of a few alternatives was most globally related to the cue. In order to manipulate the level of controlled retrieval, the possible correct alternatives varied and were either strongly (e.g., flame) or weakly (e.g., halo) related to the cue. They predicted that if the LIPFC mediates controlled semantic retrieval then weaker associate strength should elicit greater LIPFC activation even as selection demands were held constant. Wagner et al. (2001) supported this prediction, suggesting that LIPFC mediates a top-down bias signal that is recruited to the extent that semantic memory retrieval requires controlled processing. In further support of this proposal, low frequency words activate left prefrontal areas when subjects are asked to make a semantic relatedness judgment further implicating this area in controlled semantic retrieval (Chee, Hon, Caplan, Lee, & Goh, 2002). In general, the LIPFC indicates the ways that semantic memory is processed when the relative processing demands are increased. However, in cases where associative strength of becomes weaker, other more abstract semantic representations are likely to bridge the gap. The LIPFC may therefore represent an ideal location for resolving difficult category boundaries when accessing perceptually conflicting semantic features through the use of rule-based features coded indirectly in the language system.

Seminal two-stage models in categorization (Smith, Shoben, & Rips, 1974; Armstrong, Gleitman, & Gleitman, 1983) predict that when category decisions become more difficult for more borderline cases, response times for verification increases. Based on the role of prefrontal cortex in mediating controlled semantic memory retrieval and selection, the second stage of category processing is likely implemented here or in adjacent cortical regions. That is, very easy category judgments (e.g. Is a ‘dog’ a ‘mammal’?) are said to bypass the second stage completely, perhaps relying exclusively on category prototypes, perhaps in more posterior regions within inferior occipitotemporal cortex for a simple pattern match of the item and a summation of the average features within the category. However, when the decision becomes more difficult in borderline cases (e.g. Is a ‘dolphin’ a mammal?), the second stage representing the processing of abstract, or relational, properties (e.g., Does it breathe air?; Does it give birth to live young?) becomes necessary. This second stage seems to rely on rule-like semantic knowledge processed within prefrontal cortex. Recent evidence from Smith, Patalano, and colleagues (Smith, Patalano, & Jonides, 1998; Patalano, Smith, Jonides, & Koeppe, 2001) supports this prediction.

Participants were taught to use either a rule or to memorize individual exemplars when learning to visually categorize a set of novel animal-like stimuli that varied on ten dimensions (e.g., body markings and tail shape). During each stimulus block the rule was changed to prevent subjects from simply relying on memorized prototypes of the test items (see Experiment 2 in Patalano et al., 2001). The results showed that bilateral prefrontal cortical areas were activated in the rule condition but not in the exemplar memorization condition, even when time-on-task was controlled for, suggesting more directly that prefrontal areas engage in the type of feature processing required by the second stage of seminal categorization models. In this regard, prefrontal cortex would seem more generally to mediate the type of semantic processes necessary to resolve difficult category boundaries.

We have recently used a category verification task in an event-related neuroimaging design to specifically address how borderline category membership is resolved (Goldberg, Perfetti, & Schneider, 2003). Based on prior behavioral results (Goldberg, submitted) it was expected that plants (e.g., daisy, pine, etc.) would be more difficult than mammals to classify as living things since young children tend to deny that plants are alive because of their perceived lack of movement (Carey, 1985). In addition, it was expected that items with visual properties more similar to typical living things would cause greater behavioral difficulty than those instances sharing no features with the category. For example, since natural moving things (i.e., ‘cloud’ and ‘river’) share the features of ‘movement’ and ‘naturalness’ with the category *living thing*, subjects were expected to have greater difficulty denying category membership to these items, in contrast to an easier behavioral verification for artifacts, as we found in previous behavioral work (Goldberg, submitted). This added semantic difficulty was expected to drive prefrontal regions, esp. the LIPFC, based on the predicted role of these areas in representing higher-order conceptual associations, perhaps rules, and in controlling semantic retrieval demands.

As expected, the contrast between plants and mammals yielded greater activity in the LIPFC for plants. This activity in LIPFC may explain how the added difficulty for plants is resolved, perhaps via abstract properties (e.g., ‘grow’ and reproduce’). In contrast, mammals activated bilateral occipitotemporal regions implicated in visual object processing and color knowledge. This result further replicates the neuropsychological and neuroimaging evidence that shows a selective reliance on these visual regions by animal stimuli, perhaps based on a selective reliance on visual properties. As for the nonliving items, natural moving things (e.g., ‘river’ and ‘cloud’) led to greater activation in the LIPFC, the dorsolateral prefrontal cortex (DLPFC) and the anterior cingulate cortex (ACC), likely for similar reasons. Items that cause confusions in young children are more likely resolved through the acquisition and use of explicit rules, represented in the language system. In addition to replicating previous findings, these results suggest that the LIPFC is specifically activated whenever abstract properties are required. However, this finding is conflated with an increase in response time for these same items. The LIPFC may therefore represent abstract semantic properties or may more generally subserve semantic decisions as the response demands increase. The second experiment of this dissertation was specifically designed to test this prediction and is presented in Chapter 3 below.

The prefrontal cortex, based on the neuropsychological literature and more recent neuroimaging findings, appears to be intricately involved in representing and processing higher-order relational information and abstract rules. This area of cortex is the ideal candidate region for instantiating the second stage of seminal categorization models and more generally retrieving the information about a given item or category, perhaps by coordinated processing with more posterior brain regions. This region could also play a vital role in the active simulation of perceptual symbols necessary to resolve complex linguistic relations. The prefrontal cortex has been implicated in the controlled retrieval of semantic knowledge, relevant for selecting among competing alternatives, and in representing abstract rules. It is not yet clear whether any of these theories is sufficient to explain the functional significance and breadth of this region.

2. Experiment 1: Sensory brain regions represent perceptual properties

While there is much to evidence to suggest that semantic knowledge specifically relies on sensory brain regions, this conclusion is extensively debated (see for example Caramazza & Mahon, 2003). If object knowledge relies on brain regions involved in object perception, symbolic representations in language would seem to be directly related to the experiences to which they refer (Barsalou, 1999). Based on the previous neuropsychological and neuroimaging evidence, the first experiment predicted that perceptual semantic features are widely distributed throughout sensory cortical regions depending on the roles of those areas in encoding the relevant experiences with the associated objects. That is, sensory brain regions were expected to support perceptual semantic decisions in each sensory modality.

Though previous neuropsycholgical and neuroimaging studies have strongly suggested that sensory brain regions support perceptual knowledge of objects, it is not clear whether this relationship between perception and conception applies more broadly across all of the sensory modalities as a principle of knowledge representation. In addition, the use of different paradigms and stimulus classes has made clear interpretations difficult. This experiment used a common property verification task to probe knowledge of objects across all of the sensory modalities. Furthermore, by examining the activation of these regions and the LIPFC as response latency increases in modality-specific decisions, it will be possible to discriminate among multiple models of semantic processing.

The functional role of the LIPFC is typicality described as mediating semantic decisions when response times increase, perhaps to account for the necessity of selection (Thompson-Schill et al., 1999) or increasing controlled retrieval (Wagner et al., 2001). However, it is not clear to what extent this region mediates modality-general or modality-specific semantic control and whether LIPFC co-activates with sensory association cortices (see for example, Gold & Buckner, 2002). By controlling for response times within and across sensory property verifications, this experiment will examine the degree to which prefrontal and sensory regions co-activate during the resolution of semantic decisions. It addition, it will be possible to examine whether activity in LIPFC is influenced, or subdivides, by the decision modality. This experiment generally addresses how semantic information may be both highly distributed to disparate regions and integrated through processing in the LIPFC.

With increasing response times, there are at least three theoretical predictions for how modality-specific semantic regions may show increased activation either independent of or in conjunction with the LIPFC. First, feature-based accounts generally assume that semantic attributes (Martin et al., 1995) can be localized to specific regions of cortex. These areas may be sufficient to support sensory property decisions even when response times increase. In contrast, two stage models of semantic processing generally predict that modality-specific attributes will be used when sufficient to resolve semantic decisions (Smith et al., 1974). Supplementary abstract information (perhaps via inferior prefrontal cortices as addressed below in Exp. 2) may be used as necessary to resolve decisions with increasing difficulty. Lastly, controlled retrieval models (Wagner et al., 2001) would predict that LIPFC activations track with modality-specific regions as difficulty increases. This experiment aims to distinguish among these alternatives by examining how activity correlates with response times in modality-specific semantic decisions.

Feature-based accounts may predict that the spatial extent of activity in sensory regions will increase with response times during modality-specific decisions but with no other areas showing differential activation patterns. Alternatively, with increasing response times, the features used during modality-specific decisions may elicit increasing activation in the prefrontal cortex but not in modality-specific regions. This pattern of results would support a two-stage model of semantic processing in which the modality-specific regions pass information to a second stage, perhaps represented by the LIPFC, as decisions become increasingly difficult. However, if increasing response times correlate with increased activity in both prefrontal and modality-specific regions, a controlled retrieval model would be supported. That is, the LIPFC may aid in controlling processing in modality-specific regions as semantic decisions become increasingly difficult, perhaps through selecting the relevant features for the decision at-hand.

Furthermore, two-stage models and dynamic processing theories may be explained in terms of dedicated connections between the LIPFC and sensory brain regions (Gold & Buckner, 2002). These connections may drive semantic control processes or inferior prefrontal cortex may process semantic information independent of input modality. The first experiment is also designed to test these predictions: whether the LIPFC can be fractionated by semantic modality or whether this area operates as a modality-general semantic processing region. If the LIPFC represents a modality-based topography of control regions, then it should subdivide based on the modality of the property verification. Alternatively, if the LIPFC processes semantic information independent of semantic modality, then no such subdivision should be evident across the sensory-based property verifications. By examining activity in the LIPFC as semantic modality is varied, this experiment aims to investigate the top-down role of this region in controlling the activation of semantic properties in sensory brain regions.

The first experiment aimed to examine the degree to which sensory regions and left inferior prefrontal cortex (LIPFC) co-activate during the resolution of difficult semantic decisions. Specifically, this experiment is designed to test theoretical predictions regarding how semantic processing may be localized to modality-specific brain regions and rely upon a distributed semantic network that is sensitive to increases in response time. This experiment therefore examined how semantic knowledge may be both highly distributed to sensory cortical regions and accessed with increasing control depending upon task demands.

Methods

Participants

Fifteen right-handed, native American-English speakers (mean: 21.87; range: 18 – 27; 8 males) from the University of Pittsburgh community volunteered to participate in this fMRI study. Due to excessive head motion (> 3 mm) during the scanning session, the data of one female participant was not included for group analyses. The exclusion of this participant resulted in a slightly younger sample (mean: 21.86). All participants were unaware of the specific purposes of the study but were informed that the research was being conducted to examine how the human brain represents word meaning. The same participants contributed to this experiment and Experiment 2 presented below.

General Procedures

Prior to the experiment every participant completed a brief version of the property verification paradigm (described below) so they could become familiar with the task constraints including the nature of and time required for responses. In this practice version of the paradigm, items and properties were presented that did not appear during the actual experiment. The participant completed this practice version while seated at a standard computer and responses were conveyed through keyboard button presses that corresponded to the fingers required for operating the button response unit (BRU) in the MRI scanner. To mimic the noisy environment of the MRI scanner, participants heard through headphones a recording of the repetitive knocking sound made by the scanner.

All MRI scanning was conducted at the Brain Imaging Research Center of the University of Pittsburgh and Carnegie Mellon University on a 3-Tesla head-only Siemens Allegra magnet. Participants were required to lie in a supine position with their heads inside the magnet bore and viewed the stimulus display through a mirror positioned in front of their eyes that reflected the projection screen positioned behind their head. Passive restraints (e.g., compressible pillows) were used to lessen participant head motion.

The scanning session began with the collection of scout, in-plane, and volume anatomical series while functional scans corresponded to runs of the experiment. The scout images were used to position the straight-axial in-plane series (T1-weighted inversion recovery pulse sequence) at a slice prescription parallel to the inferior aspect of the temporal lobes of each participant. The 45 slices of this series provided almost full coverage of the brain for all participants. The volume anatomical scan was acquired using the Siemens MPRAGE sequence which allows for a three dimensional reconstruction of each participants brain and thus more precise anatomical localization with a resolution of approximately 1 mm^3. The in-plane structural scan served as the anatomical reference for all functional series which were collected in the same slices using a T2*-weighted echo-planar imaging (EPI) pulse sequence (TE = 30, TR = 3000, FOV = 210, slice thickness = 3.0 mm with no gap between slices, flip angle = 90, in-plane resolution = 3.125 mm^2). Functional data were collected in 4 separate runs with two successive runs devoted to each experiment. Each run in both Experiment 1 and Experiment 2 lasted for 7 minutes 42 seconds, and included 151 image acquisitions, with the order of the experiments counterbalanced across participants. The first three image acquisitions of each run were automatically discarded by the scanner and corresponded to the instruction period for the first property verification block. All verification trials corresponded to one functional acquisition (i.e., scanner TR for EPI sequence = trial length) therefore for both experiments the paradigm was a mixed event-related design with verification questions blocked. In the first experiment, the event trials correspond to easy and hard items.

During each experiment, participants were asked to determine whether an item (e.g., ‘lettuce’) possesses a given property (‘green’), a so-called property verification task. The four target properties (e.g., ‘green’ or ‘lays eggs’) of each experiment (see below) were kept constant across the experiment while the items to which participants affirmed or denied the given property were randomly varied. Property verification questions were presented in blocks of six trials with each trial sequence three seconds in length. An equal number of affirmative and negative items, 48 trials of each, were presented for each verification question and within each block of eight trials. Participants had up to 1500 milliseconds (ms) to respond to each stimulus items regarding whether or not it possessed the given property. Affirmative responses were always indicated with a button press associated with the right index finger on the BRU, while negative responses were always associated with a right middle finger button press. Stimuli were followed by the presentation of a central fixation cross (+) for 1500 ms. After the eight trials of each verification block, an instruction screen was presented for eight seconds to orient the participant to the verification question in the subsequent block. A fixation cross was presented after the introduction screen for one second to prepare the participant for the subsequent trial. Property verification questions were randomly presented six times within each functional run of the scanner for a total of twelve repetitions across the two runs devoted to each experiment. A total of 96 trials were presented for each verification question. The stimuli for both experiments are presented in Appendix C.

Design of Experiment 1

In the first experiment, four verification questions were presented (see Table 1) which varied across each sensory modality. In this regard, 'modality' refers to the sensory input (i.e., visual, auditory, tactile, or olfactory/gustatory) used to code a given semantic feature. Some modality-specific semantic regions have been previously identified and tend to rely on primary and secondary sensory areas, and so for each sensory modality, specific cortical regions were expected to be preferentially activated by perceptual properties in that modality. Based on prior research, a visual property was expected to rely on the ventral temporal cortex (Martin et al., 1995), an auditory property was expected to activate the superior temporal gyrus (James & Gauthier, 2003), and an olfactory/gustatory property was expected to drive the orbitofrontal cortex (Goldberg, Perfetti, & Schneider, 2002). As the processing of gustatory and olfactory sensory input does not seem well distinguished in the human brain, we did not separately test properties of each type. To our knowledge, no prior research has shown areas in the somatosensory cortex to be active during a tactile semantic judgment though we expected such activation for a property of this sensory modality. Therefore, we expected that the sensory brain regions associated with the perceptual encoding of each modality would also support semantic properties associated with that specific input. This experiment therefore consisted of a blocked design in which eight consecutive trials were presented within each semantic modality allowing for 24 seconds for the hemodynamic response to peak and return to baseline.

Within each verification question, affirmative and negative items were varied with regard to the time typically required to verify the given property for each item. In this regard, ‘difficulty’ refers to the average response time for subjects to verify a given property across the set of items presented for each verification question. Items within each question were classified (see again Table 1) using a median split to produce sets of ‘easy’ (i.e., typically responded to in ~700 ms) and ‘hard’ (i.e., ~800 ms) items from norming data with behavioral-only participants. Both affirmative and negative items were categorized in this way to examine the effects of increasing response time on activity in sensory brain regions and in supporting semantic retrieval mechanisms in the prefrontal cortex. Within each property block, easy and hard items were equally and randomly presented as ‘events’ with each trial associated with one volume acquisition. This fast event-related design produced a ‘jitter’ to better allow for the deconvolution of effects associated with semantic difficulty. With the verification questions presented as blocks, this mixed design allowed for statistical comparisons between the modality verifications and the event trial types.

A visuo-motor letter detection task with nonword stimuli (i.e., Does ‘contain the letter ‘i’?) was used to localize prefrontal regions specifically involved in semantic processing and as a further test of the semantic role of modality-specific sensory regions. The trial and block presentation of this control task was identical to the perceptual property verifications and was randomly interspersed throughout each run of the experiment.

Materials

For each property verification, norming experiments with behavioral-only participants (see Appendix B) were used to: 1) ensure that item-property relations within each question relied on perceptual knowledge to the same degree; 2) identify clearly affirmative and negative items and items that vary according to ‘difficulty’; and 3) match the set of items within each question for possibly confounding lexical factors. One set of behavioral-only participants rated the degree to which they knew the relationship between the given item and a property. Based on a five-point scale, participants indicated whether they knew the relationship from having directly experienced it (‘sensory’) or having read about it or been taught it (‘abstract’). The modality-specific questions used in this experiment were matched on this scale. This method was also used to identify the sensory and abstract questions used in Experiment 2. Another set of behavioral-only participants performed the property verification decision on a superset of properties and items used in the experiment. Their accuracy data was then used to segregate affirmative and negative items while eliminating any item for which the property relationship was not clear. That is, any any item for which the property was affirmed or denied between 40% and 60 of the time was not included in the stimulus set (see <u>Response</u> in Appendix C). In contrast, their response time data was used to classify easy and hard instances within the affirmative and negative item sets of each verification question. Hard and easy items within each modality were separated on average by about 100 ms (see <u>RT</u> in Appendix C). Affirmative and negative items within each modality were drawn from the same superordinate categories. Stimuli in the letter detection blocks consisted of pronounceable nonwords constructed from rearranged letters of words used in the sensory property verifications.

Easy and hard trials within each modality, as well as all items across all property verification questions, were matched for (see Appendix B) letter length (Avg. = 5.96), phonological length (Avg. = 1.91), and lexical familiarity (Avg. = 0.29).

Statistical Analyses

Neuroimaging data was preprocessed and analyzed using the BrainVoyager 2000 (BrainInnovation, Inc.) software. Preprocessing steps included: 6-parameter, 2D motion-correction, slice scan time correction using sinc interpolation, voxel-wise linear detrending, and spatial smoothing with a 8 mm full-width half-maximum (FWHM) Gaussian kernel. These processing steps are necessary to account for intersession and intersubject variance including scanner artifacts, head motion, and neuroanatomical differences. To allow for group composite results, spatial normalization was performed using the standard 9-parameter landmark method of Talairach and Tournoux (1988). For analyses, a general linear model was defined for each subject that included five regressors which modeled the BOLD response to each 24 second block period (4 trials) associated with the presentation of each property verification block and the control condition. Each regressor was convolved with a standard gamma model of the hemodynamic impulse-response function (Boynton, Engel, Glover, Heeger, 1996). The resulting statistical maps from the contrasts described below were thresholded for significance and cluster size (>= 5 contiguous voxels). Any resulting region-of-interest (ROI) analyses transformed the arbitrary values of the neuroimaging data to a more informative measure of percent signal change. This latter measure assumed the run average as the baseline activity and computed the variation of each data point in each run from this baseline.

Contrasts in the neuroimaging data were made within and between the verification questions. Conjunction analyses of all pair-wise comparisons with the modality of interest were used to localize the activity in distinct sensory cortical regions predicted for each modality-specific decision. For example, to localize regions specific to the auditory property verification, a conjunction analysis of all pair-wise contrasts (i.e., Loud – Green, Loud – Soft, Loud – Sweet, and Loud – Control) revealed the overlap map of all voxels at significance for each comparison. This test is the most conservative for ensuring that only those areas activated by the predicted sensory modality, but not the others, will reach threshold in the statistical maps. Regions-of-interest were determined a priori using landmark coordinates based on previous research findings and known locations of primary sensory regions. Activity in these regions for each of the property verification decisions was then analyzed using single sample t-tests to examine the probability against chance of a significant increase from baseline activity. Candidate semantic regions in the prefrontal cortex were localized by contrasting all semantic verification blocks with the nonword letter detection task, and overlapping the resulting statistical map with one from the direct contrast between hard and easy trials. To examine the effects of 'difficulty' on activity in modality-specific regions and the prefrontal cortex, paired-sample t-tests were used to examine activity differences between easy and hard trials in regions-of-interest defined through the above comparisons. The effect of response time on activity in these regions was then examined to test the theoretical predictions of feature-based, two-stage, and dynamic processing theories. Activity in prefrontal regions was also examined for modality-specific effects between areas based upon possible dedicated connections to posterior brain regions. Behavioral data was examined using paired-sample t-tests to analyze significant effects of modality, based on response time differences between property verification decisions, and difficulty, based on response time differences between all easy and hard trials within each modality and across the sensory modalities. Single factor ANOVA tests were used to examine the main effect of semantic difficulty between easy and hard trial types.

Behavioral Results

The behavioral results show few differences between sensory modalities but do support the difficulty manipulation between easy and hard items within each modality. Among the sensory verifications, the visual property was responded to slightly more quickly, by about 30 milliseconds (ms), than the auditory and tactile decisions (see Figure 1; t's > 2.4, p's < .05). All other response time differences between the sensory modalities were not significant. Participants were slightly more accurate for the auditory verification, by about 5 %, than for the gustatory property (t = 3.0, p < .05). No other accuracy differences between the modalities were found. The letter detection control was responded to more than 160 ms faster (t's > 9.87, p's < .001) and 6% more accurately (t's > 5.25, p's < .001) than all of the sensory decisions. Overall, the accuracy of neuroimaging participants was 90%, indicating that the responses for items greatly matched the responses given by the participants in the behavioral norming studies. Here the neuroimaging participants were accurate on: 90% of the trials for the visual decision, 92% of auditory trials, 89% of the tactile items, and 87% of the items for the gustatory verification, while accuracy for the control condition was 97%.

As expected, based on the manipulation of item response time, there were main effects of difficulty across the sensory decisions. Easy items were responded to on average about 80 ms faster than hard items (F = 19.36, p <.001) and about 8% more accurately (F = 33.72, p <.001). Similarly as expected, within all of the perceptual properties, easy and hard trials, as defined a priori from the norming experiments, differed significantly between 65 to 100 ms (t's > 3.91, p < .001; for details see Figure 2). When using the actual data from the neuroimaging participants, the average response time difference between easy and hard items increases to about 275 ms (F =

212.09, $p < .001$) and within the sensory modalities this difference ranges from about 260 to 295 ms (t's > 15.16, p's < .001; see Figure 3).

Neuroimaging Results

The cortical distribution of perceptual property knowledge to sensory brain regions

Sensory brain regions were activated by the verification of perceptual properties across the four sensory modalities at a significance threshold of $p < .01$ (uncorrected). The sensory modality directed examined in each property verification question was specifically associated with the most robust increases in activation of the predicted sensory brain regions (see Table 2). This activity did not vary across affirmative and negative items. The visual property verification was specifically associated with increased activity in the left ventral temporal cortex (Figure 4a; x: -59, y: -46, z: -6), a region previously shown to be involved in the representation of object color knowledge and more generally involved in processing features used in visual object recognition. This ventral temporal activity differs significantly ($t = 3.81$, $p <.01$) by .15% from the average activity over the course of the runs while none of the other property decisions show such selective increases (Figure 4b). In addition, this verification decision was also associated with increased activity in superior parietal regions implicated in judgments of object form (Oliver & Thompson-Schill, 2003). By comparison semantic property decisions regarding whether items 'sound loud' was specifically associated with increased activity centered on the left superior temporal sulcus (STS) just inferior to primary auditory cortex (Figure 5a). The STS has been implicated in the integration of visual and auditory knowledge and in retrieving sounds associated with objects. Activity in this region increased significantly by about .075% ($t = 2.01$,

$p < .05$) for the auditory property verification in contrast to the average baseline and decreased activity for the other verifications (Figure 5b).

We expected the somatosensory cortex to show the greatest relative increase in activity for the tactile decision based on the predicted relationship between sensory encoding and perceptual object knowledge. Somewhat, but not too, surprisingly, verification of whether items ‘feel soft’ was associated with increased activity in left lateralized regions including the somatosensory cortex but also in the primary motor cortex and pre-motor cortex (Figure 6a). The tactile property decision led to activation increases in the somatosensory cortex of the left post-central gyrus by about .1% signal change over and above average activity ($t = 3.69$, $p < .01$) in contrast to the other sensory modalities (Figure 6b). The primary motor cortex in the left pre-central gyrus (Figure 6c) and a pre-motor region in the left prefrontal cortex (Figure 6d) showed similar patterns of activation of about .1% specific for the tactile property ($t = 3.83$, $p < .01$) but not for the other verification decisions. This robust activity throughout left-lateralized motor and somatosensory regions would seem to indicate the off-line simulation of tactile properties through a gestural movement toward the associated object.

As predicted, the verification of whether an item ‘tastes sweet’ was associated with specific increased activity in the left orbitofrontal cortex (Figure 7a). The gustatory modality (Figure 7b) shows a significant increase of .15% relative to the run average baseline ($t = 3.29$, $p < .01$), in contrast to the other sensory modalities. This result provides direct support for our previous interpretation of the role of this region in the representation of gustatory properties of categories that require it (e.g., fruit), though here the activity is less robust in terms of spatial extent and is restricted to the left hemisphere.

The role of sensory and prefrontal brain regions in supporting extended semantic processing

The relative difficulty of verifying perceptual properties was varied in order to test the role of sensory and prefrontal brain regions in supporting semantic processing. Contrary to expectations, the a priori classification of easy and hard trials, while associated with behavioral effects within the sensory modalities, was not associated with increased activity in the modality-specific regions discussed above for each of the modality-specific verifications. Even when the actual response times of the neuroimaging participants were used to classify easy and hard trials, which reveal a four times increase in effect size for response times within the sensory modalities, no significant effects were found in the set of sensory brain regions specifically identified for each perceptual decision. In this regard, modality-specific sensory regions would not seem to interact with the difficulty of the property verification. These regions would instead seem to contribute a threshold level of activity for the given perceptual property decisions.

On the other hand, prefrontal regions typically ascribed to modality-general aspects of cognitive control were specifically involved in more difficult sensory verifications and depending on the modality (see Table 2). A region in the LIPFC (Figure 8) was associated with increased activity for the gustatory and tactile properties (t's > 2.17, $p < .05$) but irrespective of item difficulty. That is, within these modalities, this region was not associated with increased activity for the items requiring a longer latency, as defined from the behavioral data of the neuroimaging participants. In contrast, the visual and auditory property verifications did not seem to specifically drive this region over and above the average baseline activity. However, within the auditory decision there was a trend toward significance ($t = 1.47$, $p < .09$) with increasing activity for the harder items in contrast to the less difficult items.

A more dorsolateral aspect of the left prefrontal cortex (DLPFC) appears to be more sensitive to both the probed semantic modality and response latency. This region (Figure 9) is most responsive to the tactile modality with activity increasing about .1% over the average baseline (t = 3.15, p < .01). While there were no differences between easy and hard trials for this verification decision, there was significantly increased activity for the gustatory modality. The more difficult verifications of whether an item 'tastes sweet', as defined from the behavioral data of the neuroimaging participants, were associated with increased activity of about .04% above activity for the easier trials (t = 1.91, p < .05). This specific increase in activity for harder gustatory verifications is reflected in a .07% increase above baseline (t = 2.28, p < .05) whereas the activity for easy trials did not differ from the run average. The DLPFC therefore appears selectively responsive to both the probed semantic modality and the difficulty with which the semantic properties of items are resolved.

While these regions show increased activity for specific semantic modalities and the difficulty of semantic decisions, from this data there does not appear to be an organizing principle for the controlled retrieval of semantic knowledge in the prefrontal cortex. In this regard, the resolution of semantic decisions may rely on modality-specific prefrontal mechanisms and modality-general processes alluded to in other cognitive domains. In support of this interpretation, a region in the anterior cingulate cortex (ACC) showed increased activity for the more difficult property verifications across three of the four sensory modalities (Figure 9). Even as only the gustatory verification showed significantly sustained activity from baseline for all trials (t = 3.47, p <.01), the more difficult items in the auditory, tactile, as well as the gustatory modalities led to significantly increased activity in contrast to the easier trials (t's > 1.90, p's < .05). This involvement of the ACC would to seem to indicate its more generally described role in executive control mechanisms, specifically here in the retrieval of semantic knowledge, either through the detection of cognitive conflicts (MacDonald et al., 2000) or increasing cognitive demands (Schneider & Chein, 2003).

Discussion

This experiment tested whether sensory brain regions support the verification of perceptual semantic properties of object knowledge, as predicted by perceptual symbol systems (Barsalou, 1999) and feature-based models of semantic knowledge (e.g., Martin & Chao, 2001). Across the visual, tactile, gustatory, and auditory modalities, perceptual property verifications of word items were associated with increased activity in brain regions typically ascribed to the sensory encoding of each modality. The left ventral temporal cortex typically associated with visual object recognition, but also implicated in knowledge of object colors, was specifically activated when participants verified whether an item looks green. The left superior temporal sulcus, a region involved in auditory processing, was specifically activated by the verification of whether an item sounds loud. Interestingly, the tactile verification specifically activated not only the somatosensory cortex in the left post-central gyrus, but also activated the left primary and pre-motor cortices. This result suggests that in accessing knowledge for whether an item feels soft, participants simulate the action with the specific object in retrieving the answer. Finally, the verification of whether an item tastes sweet was associated with increased activity in the left orbitofrontal cortex. This region is one we have previously identified based on its distinctive activation for food items. Based on these results, perceptual knowledge of objects and sensory experiences of them appear to rely on a common neural substrate. This relationship between knowledge of objects and sensory experiences of them suggests how world knowledge more generally arises from our interactions with the world. In this regard, the encoding of sensory experiences would seem to ground later perceptual simulations independent of the objects that caused them.

Perceptual symbol systems, as outlined by Barsalou (1999), are best suited to describing how this relationship unfolds through instances of objects and use of selective attentional mechanisms to home in on, and elaborate, specific attributes of those objects. The resulting symbolic representation of the object (e.g., a word) gradually comes to impart, or refer to, simulations of the experienced object. According to this theory, the verification of perceptual properties would therefore involve sensory mechanisms involved in experiencing the tested objects. In supporting this prediction, the present results suggest that such explicit control of prior sensory episodes is involved in the recollection of knowledge for properties like ‘green’, ‘loud’, ‘soft’, and ‘sweet’ across the spectrum of sensory modalities. Other properties within each of these modalities would similarly be expected to activate the associated sensory brain regions.

These findings provide further support for the view that semantic knowledge is widely distributed throughout the cerebral cortex depending on the specific types of features involved (e.g., Martion & Chai, 2001), rather than through more general domain-based differences (Caramazza & Shelton, 1998) or even undifferentiated at the content level (Tyler & Moss, 2001). The topography of semantic memory would therefore seem to depend on the specific properties being accessed either individually or through more superordinate relations such as those within object categories. In this regard, these results provide converging support for the perceptual / functional hypothesis originally derived from category-specific effects found in some neuropsychological patients. The selective loss or preservation of object categories would seem to result from the differential reliance on specific types of properties across classes of items.

Categories that rely more on visual properties (e.g., animals) are impaired with damage visual brain areas whereas these items are spared when lesions do not disrupt these regions. The prediction that follows from the present results is that category knowledge that preferentially relies on other attribute modalities (e.g., wines or textiles) will be specifically impaired with damage to the supporting sensory brain regions. Similarly, while certain categories (e.g., animals) may remain intact after localized damage to one modality (e.g., vision), properties of items are expected to be specifically impaired if damage occurs to areas supporting other sensory knowledge of items (e.g., barks or meows?).

While these results suggest that sensory brain regions are intricately tied to the knowledge of perceptual properties, it is not clear whether these regions are automatically activated with the meaning of a word. That is, the retrieval demands of the verification task may have induced the equivalent of an imagery strategy in each of the tested sensory modalities. Alternatively, the retrieval of semantic knowledge associated with perceptual properties may always lead to the activation of sensory brain regions. In support of this latter interpretation, Hauk and colleagues (Hauk et al., 2004) have recently found that passive reading of motor-based words induces activity in the associated regions of the primary motor cortex. Participants in the MRI were initially asked to move their hands, feet, and tongue to map regions in motor cortex specifically involved in each of these movements given the somatotopic organization of this area of cortex. In the experiment, words like pick, kick, and lick were presented within a random stream of other words that did not have a motor component. Hauk et al. found that words like ‘pick’, that include a manual component in their meaning, were associated with increased activity not only in the motor cortex but to the greatest extent in the region associated with actual movements of the hand. Similarly, words like ‘kick’ and ‘lick’ were associated with the most robust activity increases in the regions mapped to movements of the feet and mouth respectively. These results suggest that sensory and motor brain regions are automatically activated by the meaning of words in passive reading rather than through particular task demands.

The manipulation of item difficulty was used to examine the role of sensory and prefrontal brain regions in supporting extended semantic processing. Increased response time in modality-specific semantic decisions was not associated with increased activity in the sensory brain areas. This overall finding argues against strictly feature-based and controlled retrieval models of semantic processing. Rather, prefrontal brain regions appear to selectively activate as a second stage in semantic verifications based on the tested modality and the difficulty of individual items.

A region in the LIPFC was activated by the tactile and gustatory verifications, but within these modalities there were no differences between easy and hard trials. The visual and auditory modalities showed no selective activation in this region though the auditory modality exhibited a trend of increasing activity for the more difficult verifications of properties relating to object sounds. This result suggests that different regions of the prefrontal cortex support retrieval of semantic knowledge depending on the perceptual properties involved and the necessity of controlled processing mechanisms. In this regard, the controlled retrieval of semantic knowledge may depend on cortico-cortical pathways between prefrontal and specific sensory brain regions.

This interpretation is supported by the activity patterns found in the DLPFC. This region was also specifically activated by the tactile verification without regard to the difficulty of the items. However, the harder gustatory decisions were associated with increased activity in this region in contrast to baseline activity for the easier verifications within this modality. In this regard, the DLPFC appears sensitive to both the modality of the semantic decision and the time in which that decision is resolved. This specific activation of both the LIPFC and the DLPFC for the tactile and gustatory properties is not evident from the behavioral effects. It is not clear why these modalities should lead to an increased reliance on prefrontal control mechanisms. These prefrontal regions may be more involved depending on the role of prior sensory experiences (e.g., direct taste or touch information) in mediating the item-property relationship, whereas other relationships while still based on sensory mechanisms are more likely inferred or observed through indirect means (e.g., a cannon firing in a movie). This hypothesis requires further testing to delineate the roles of direct and indirect sensory experiences on the representation of perceptual knowledge.

The activation of the ACC fits with the interpretation of prefrontal regions as the second stage of semantic processing with constraints from the sensory modality of the decision and the difficulty with which it is resolved. More difficult items in the auditory, tactile, and gustatory modalities, led to increased activity in this region in contrast to the easier items within each of these modalities. This result suggests that other cognitive processes, perhaps based on the mediation of conflicting responses, are necessary for the more difficult property verifications. This interpretation of the activity in the ACC is consistent with prior models of the role of the prefrontal cortex in cognitive control mechanisms and extends these theories into the semantic domain. The lack of similar increases in modality specific sensory regions argues against strictly feature-based or controlled retrieval models of semantic processing.

Sensory brain regions would appear to support semantic decisions that rely on the perceptual properties. These regions would seem to underlie modality-specific semantic decisions regardless of the difficulty of the decision and instead contribute a base level of activity from which the verification can be made. This activity may be akin to the contribution of whether an item possesses a feature in the modality with more general mechanisms necessary for resolving whether the item meets the criteria of the question. For example, auditory cortex may represent whether an item makes a sound while prefrontal regions resolve the scalar judgment of whether an item sounds 'loud'. In this regard, the sensory regions may automatically activate for modality-based decisions, while prefrontal control mechanisms use this information to determine whether the item meets the semantic criterion as indicated by other prefrontal regions in service of selections of goal-based actions (Matsumoto et al., 2003). This second stage would seem more necessary for more difficult decision, while very easy verifications could simply rely on the base level of activity in the sensory regions. A design more sensitive to individual trials would need to test this prediction against possible alternatives.

These results argue against strictly feature-based accounts (e.g., Martin & Chao, 2001) and controlled retrieval theories (e.g., Wagner et al, 2001) that specify the LIPFC as the locus of extended semantic processing. Previous results identifying the LIPFC in semantic retrieval process may be biased by the verbal, and thus more abstract, nature of the stimulus materials, as further indicated by Experiment 2 below. The prefrontal cortex more generally may be involved in difficult semantic decisions depending on the sensory modality involved and the necessity of more abstract semantic representations. Sensory brain regions may therefore contribute the necessary content required in a given semantic decision while prefrontal brain regions are selectively activated to mediate control of that content. These neuroimaging results, and their convergence with category-specific deficits, suggest that object categories, and word meanings more generally, are instantiated through neural systems used to encode the relevant sensory information. In this regard this experiment supports a dynamic feature-based account for the organization and interaction of semantic knowledge. Sensory brain regions appear to support the representation of perceptual properties associated with objects while prefrontal mechanisms may operate over a base level of activity when necessary to resolve the semantic decision. In this regard, the prefrontal cortex may contribute top-down bias signals that control processing in posterior cortical regions (Miller & Cohen, 2001). This is a point that we will return to in the general discussion below.

One problem with this sensory-based perspective is that the entirety of semantic knowledge cannot be reduced to perceptual terms. For instance, as described earlier, children are often confused by their sensory experiences with objects and must learn to inhibit prepotent responses based on perceptual cues to learn theoretical knowledge (e.g., Carey, 1985). Overcoming this difficulty involves acquiring theory-based features to describe conceptual events and situations that cannot be directly experienced. In this regard, perceptually-biased feature-based models do not specify the properties used in more abstract associations, such as when prefrontal cortices are shown to be involved in rule-based processing associated with perceptual categorization (Miller, Freedman, & Wallis, 2002). The second experiment aimed to examine the content by which some prefrontal brain regions operate over semantic knowledge. That is, even as the first experiment indicates the role of dissociable prefrontal regions in the retrieval of semantic knowledge, it is unclear what the grist is for this processing mill.

3. Experiment 2: Prefrontal regions represent abstract semantic knowledge

The functional role of left inferior prefrontal cortex (LIPFC) is typicality described as mediating semantic decisions when response times increase (Roskies et al., 2001). However, semantic decision experiments that focus on the LIPFC often confound this extended processing with the greater abstractness of the stimulus materials (Wagner et al., 2001). It is not clear whether the extended processing necessary for some semantic decisions is resolved through abstract information stored in the LIPFC. Since neurons in the primate analogue of this region are seen to represent abstract rules (see for example, Wallis et al., 2001), it is necessary to distinguish in humans how this region may either represent abstract knowledge or is more generally activated by demanding semantic decisions.

The second experiment used property verification decisions to dissociate semantic difficulty and degree of item-property abstractness on activity in the LIPFC. Participants were presented with words and asked to verify whether a property belongs to the meaning of the item. Property verification questions were controlled for how long, on average, it takes subjects to respond to all items (‘difficulty’) and the degree to which an item possesses the property through perceptual experience or learned indirect knowledge (‘abstractness’). By examining activity in LIPFC as either difficulty or abstractness is held constant, this experiment aims to clarify the role of this region in semantic processing. Based on these variables, it may be possible to dissociate anterior and posterior aspects of the LIPFC based on the control of semantic knowledge and the abstract content of some decisions. This result would complement the prefrontal subdivisions found in Experiment 1.

The second experiment aims to clarify the role of the LIPFC in processing abstract semantic information. This experiment examined how cortical activity in the LIPFC varies depending on task difficulty and/or the degree to which semantic information is separated from perceptual experience and rather learned indirectly through the language. The goal of this experiment was to disambiguate the effects of semantic difficulty and the abstractness of semantic information on activity in the LIPFC. Abstract semantic properties may specifically rely on the LIPFC and may more generally explain the activation of this region during demanding semantic tasks. If the LIPFC is seen to support abstract semantic processing irrespective of response times, this evidence would further implicate this region in two-stage models of categorization.

Methods

Participants

The same participants contributed to Experiment 1 and this experiment. Group analyses therefore included fourteen right-handed, native American-English speakers (mean: 21.86; range: 18 – 27; 8 males) from the University of Pittsburgh community.

Design of Experiment 2

In the second experiment, four verification questions were presented (see Table 3) which varied across item-property abstractness and response difficulty. Accordingly, 'abstractness' refers to the ratings given for item-property relations depending upon whether the relationship relies on directly perceived experience ('sensory') or indirectly learned knowledge ('abstract').

Verifications were therefore classified, based on the behavioral ratings (scaled 1 to 5 with 5 being very abstract), as either sensory (e.g., < 3) or abstract (e.g., > 3). These ratings were expected to vary for each item within a property. By averaging across all items, some property decisions were more sensory (e.g., four legs?) and other were more abstract (e.g., can be trained?). In contrast, verification questions were also varied with regard to the response time typically required to verify the given property for each item. Therefore 'difficulty' refers to the average response time for subjects to verify a given property across the set of items presented for each verification question. However, in this experiment, rather than classify items as either 'easy' or 'hard', the verification questions more generally were classified in this way based on the average response time across all items used within the question. That is, averaging across the set of items within each verification question produced decisions that were either 'easy' (i.e., average response time for the question equal to ~ 700 ms) or 'hard' (i.e., ~ 800 ms). As shown in Table 3, this experiment therefore contrasted four questions based on whether they varied according to abstractness or difficulty. By controlling the materials in this way, this experiment independently examined the effects of 'abstractness' and 'difficulty' on activity in the LIPFC depending on the verification. This mixed design allowed for contrasts between the blocks of the verification decisions while segregating the event-related data, including one volume acquisition for each trial, for each incorrect response.

A visuo-motor letter detection task (e.g., Does X have contain the letter 'e') was used to localize the LIPFC and other prefrontal regions specifically involved in semantic processing and the representation of abstract semantic knowledge. The trial and block presentation of this control task was identical to the property verification questions and was randomly interspersed throughout each run of the experiment.

Materials

For each property verification, norming experiments with behavioral-only subjects (see Appendix B) were used to: 1) identify verification questions that vary according to 'abstractness'; 2) identify verification questions that vary according to 'difficulty' and to identify clear affirmative and negative items within each question; and 3) match each question for possibly confounding lexical factors. One set of behavioral-only participants rated the degree to which they knew the relationship between the given item and a property. Based on a five-point scale, participants indicated whether they knew the relationship from having directly experienced it ('sensory') or having read about it or been taught it ('abstract'). This method was also used to match the sensory verifications used in Experiment 1. Another set of behavioral-only participants performed the property verification decision on a superset of properties and items used in this experiment. Their accuracy data was used to segregate affirmative and negative items while eliminating any item for which the property relationship was not clear. The response times of the remaining items for verifying each property was then averaged to classify each potential question as either 'easy' or 'hard'. Therefore, the behavioral data allowed for the selection of verification questions, through the items within them, to maximize two of the four possible extremes with regard to 'difficulty' and 'abstractness' (see again Table 3): 1) easy and sensory (Avg. RT = 671.55 ms; Abstractness = 1.71); 2) hard and sensory (Avg. RT = 769.46 ms; Abstractness = 1.72); 3) easy and abstract (Avg. RT = 710.63 ms; Abstractness = 3.54); 4) hard and abstract (Avg. RT = 822.81 ms; Abstractness = 3.47). Importantly, all affirmative and negative items for the four property decisions were drawn from the same superordinate category (i.e., organisms) and no items were repeated. The four property verification questions were also matched for letter length (Avg. = 5.78), phonological length (Avg. = 1.83), and lexical familiarity (Avg = .34).

Stimuli in the letter detection control consisted of pronounceable nonwords constructed from rearranged letters of words used in the property verifications. This localizer task was used to isolate prefrontal regions specifically involved in semantic decisions. This task differed from the one used in the first experiment in terms of the letter detected (i.e., ‘e’ versus ‘i’) and the nonwords used.

Statistical Analyses

Behavioral data was examined using two factor ANOVA tests to analyze significant effects of abstractness and difficulty, as well their interaction. Accuracy data was used to specifically examine response times and neuroimaging data for correct trials only. The neuroimaging data was preprocessed and analyzed using procedures identical to those of the first experiment. Contrasts in the neuroimaging data were made between the verification questions based on activity differences in regions-of-interest (ROI’s). Candidate ROI’s in the prefrontal cortex were localized by contrasting all semantic verification blocks with the nonword letter detection task in a weighted GLM. To examine the independent contributions of abstractness and difficulty, and their interaction, on activity in these prefrontal ROI’s, two factor ANOVA tests were conducted on the percent change neuroimaging data.

Behavioral Results

The behavioral results show main effects of response time and accuracy for both abstractness and difficulty but no interaction. As shown in Figure 11, there were significant effects of abstractness and difficulty on overall accuracy to the property verifications. The two abstract questions were responded to on average about 11% less accurately than the sensory verifications (F= 52.62, p < .001). Similarly, responses for the two hard questions were generally about 7% less accurate than for the easy questions (F = 19.28, p < .001). The interaction of abstractness and difficulty on accuracy was not significant. The response time data also show main effects for abstractness and accuracy but no interaction (Figure 12). Participants responded about 110 ms slower for the two abstract questions than for the two sensory questions (F = 22.22, p < .001) while responding about 60 ms more slowly for the hard than for the easy verification decisions (F = 6.29, p < .01). Participants encountered much less difficulty on the nonword control task in responding, on average, in about 750 ms with an accuracy of approximately 97%.

Neuroimaging Results

Abstract property verification specifically activates the left inferior prefrontal cortex

By contrasting the property verifications with the nonword control task, candidate semantic regions were localized in the left inferior prefrontal cortex to independently investigate the effects of abstractness and difficulty on activity in this region. This comparison yielded two ROI's in the LIPFC, one more medial and the other more lateral, specifically in the most ventral aspect of Broadmann Area (BA) 47 (see Table 4). The activity in both of these regions show main effects of abstractness but no main effect of difficulty however with noteworthy differences based on the interaction of these variables. The abstract property verifications produce the most activation in the medial aspect (Figure 13; x: -31, y: 35, z: 0), by about .1% over the perceptual properties (F = 7.51, p < .01) at a Bonferroni corrected threshold of p < .01. There is no significant effect of response time nor is there an interaction of abstractness and difficulty on activity in this region. However, the more lateral aspect (Figure 14), at a slightly more lenient Bonferroni corrected threshold of p < .05, shows both a main effect of abstractness (F = 55.09, p

< .001) and a significant interaction (F = 8.87, p < .01) of abstractness with difficulty. The abstract properties lead to significantly increased activity in this region by about .3% over the perceptual verifications. Interestingly, the interaction of abstractness and response time is driven not by the more difficult abstract property, but rather by the easier decision. The verification of whether an item 'can be trained' leads to average increased activity of about .17% (t = 4.11, p < .01) over and above verifications of whether organisms lay eggs. This result may reflect how abstract properties depend on certain other facts in semantic knowledge and could result from subtle differences between these particular properties. In contrast there is no significant difference in activity between the two sensory verifications. The combined activity across these two regions suggests that rather than operate over more general semantic processing, the LIPFC appears to be specifically involved in processing abstract semantic properties. However, this area of the prefrontal cortex my not be solely responsible for the representation of this type of semantic knowledge.

The prefrontal cortex more generally processes abstract semantic knowledge

The contrast between the property decisions and the letter detection control also yielded two additional regions in more dorsal aspects of the left prefrontal cortex, specifically in the middle frontal gyrus with one more anterior and the other more posterior at p < .01 (Bonferroni correction). These regions similarly showed significantly increased activity for the abstract properties in contrast to the sensory-based decisions. However, as with the more inferior regions, these areas also show subtle differences in terms of the interaction between abstractness and question difficulty. The more anterior region responds very similarly to the lateral aspect of the LIPFC. This region (see Figure 15) shows both a significant main effect of abstractness and a significant interaction of abstractness with response time. The abstract verifications, on average, led to a significant increase of .2% in contrast to the perceptual verifications (F = 20.20, p < .001). As in the medial LIPFC area, the significant interaction (F = 6.02, p < .05) between abstractness and response time in this region appears to be largely driven by an increased response to the easier abstract question. When participants verified whether an organism can be trained, activity in this region significantly increased, by about .12%, over the more difficult 'lays eggs' verification (t = 2.40, p < .05). In contrast to this more anterior region, the more posterior dorsolateral aspect of the prefrontal cortex (x: -44, y: 21, z: 32) was responsive to the abstractness of the verification questions, but showed no effects of response time. The abstract verifications were associated with a significant increase (F = 14.53, p < .001) of .1% over the perceptual decisions. This DLPFC region showed no effects of behavioral response time on activity.

Discussion

This experiment examined the role of the prefrontal cortex in supporting abstract semantic decisions as response times increase. The results suggest that the LIPFC is involved in representing and processing abstract semantic properties in contrast to perceptual properties. More dorsal prefrontal regions also appear to be associated with semantic decisions as abstractness increases. That is, rather than indicate that one region in the prefrontal cortex is a simple repository for abstract properties, these results suggest that the prefrontal cortex more generally supports the representation and processing of this type of semantic knowledge. Disparate and neighboring prefrontal regions may show specializations based on a number of factors indicating modality-specific and domain-general mechanisms. The representation and controlled processing of semantic knowledge in the prefrontal cortex may depend not only on the abstractness and difficulty of the decisions but also, as indicated in Experiment 1, the sensory modality from which perceptual properties are derived.

Across the tested verification questions, abstract item-property relations specifically led to increased activity in the left inferior and middle prefrontal cortex. This finding suggests that previous studies that have found prefrontal involvement in semantic tasks may not have sufficiently controlled for the abstractness of the stimulus materials or the specific decisions required. That is, the prefrontal cortex may become increasingly necessary as semantic decisions become more abstract, or the associations weaker (Wagner et al., 2001), rather than simply more difficult. This experiment controlled for the abstractness of semantic properties while all of the items represented concrete objects. In this way, the role of prefrontal cortex would seem to be more directly associated with the abstractness of the retrieved semantic information rather than other stimulus factors. It is important to note that these results should be qualified by pointing out that only two abstract properties were tested. The prediction that follows though is if semantic decisions are controlled for their abstractness, perhaps as suggested from the behavioral norming studies, prefrontal regions become increasingly necessary irrespective of the property. Yet this norming method raises the question of what exactly it means for semantic knowledge to be abstract and whether the chosen method is an adequate approach.

In the behavioral norming experiments, participants were asked to determine the degree to which they knew that a given property belongs to a specific object. They rated whether they knew this relationship from direct sensory experiences with the objects or through more abstract means, like having read about or been taught the relationship. In this regard the abstractness of the item-property relationship was defined through relying either on perceptual or verbal knowledge. The instructions for this behavioral norming were slightly modified from previous experiments in which participants are asked to rate on a 7-point scale the concreteness of particular lexical items (see for example Toglia & Battig, 1978). For example, in these ratings an item like ‘table’ is rated as very concrete while an item like ‘justice’ is rated as very abstract. Since the aim of the second experiment was to examine the abstractness of item-property relations, the instructions for the norming study (see Appendix B) were altered to redirect participants to their knowledge of properties that belong to a given item. The instructions required participants to rate whether they knew that concrete items had a given property through direct sensory experiences with objects with or more indirectly through the language, such as having been taught or read about the relationship. The important point is that this rating was based on a sliding scale and so differences between properties are likely more a matter of degree than of kind, especially since all of the items represented concrete organisms, and could be directly experienced. Yet this norming dimension may be one of a few possibilities for how semantic abstractness may be judged and manipulated.

Barsalou (2003) has recently identified at least six possible senses of abstraction and the ways in which it can be described. Category knowledge (i) refers to the abstraction away from the experience of individual exemplars to form a lexical representation independent of them. In this regard, ‘chair’ refers to the set of items in the world that tend to share a certain physical structure with a given intended usage. Two additional senses of abstraction refer to either the behavioral ability to generalize across such category members (ii) or the sense in which a given abstraction, such as ‘chair’, can be stored as summary representation in memory (iii). The difficulty with this third sense is in the type of format this stored representation takes and theories differ according to whether this summary representation is a prototype of shared features (e.g., Rosch & Mervis, 1975), an explicit feature (e.g., Smith, Shoben, & Rips, 1974), or even a set of connection weights (e.g., Rogers & McClelland, 2004). Barsalou (2003) also identifies three additional senses which abstraction refers to: (iv) schematic representations less specific than individual exemplars that (e.g., the geon theory of object recognition - Biederman, [1987]), (v) flexible representations in which a given abstraction can be applied independent of particular relations to things in the world, and finally (vi) abstract concepts, as studied in the traditional norming studies, in which lexical items can refer to concrete objects, like 'tiger', or more abstract ideas, like 'justice'. Barsalou (2003) suggests that the third sense, rather than depend on static summary representations, arises from how such relationships develop from dynamic simulations of objects and their constituent properties.

According then to these senses, abstract properties of concrete items, of the type examined in this experiment, may arise from category relations (e.g., mammal), the behavioral tendency to form such abstractions, or stored representations in memory that may or may not be strictly categorical (e.g., things that fly). In addition, by relying on norming studies that draw on the distinction between abstract and concrete word meanings, this experiment more generally illustrates how the particular representations, or simulations of them, may change based on the semantic content involved. For instance, while the four feet of a coyote may be directly experienced through the visual system, that a coyote can or can not be trained would seem to require knowledge independent of what it looks like or how it behaves. In this regard, the language system may provide more indirect knowledge associated with particular items but not necessarily reliant on any particular category relation or modality.

Abstract properties may be also bound more tightly to the context in which they are generated and requiring particular events or circumstances in a given situation (Wiemer-Hastings et al., 2001). These contextual relations may be more likely to be mediated through linguistic knowledge rather than stored as language-specific representations, as indicated by the norming studies for abstract properties. So, for instance, ‘lays eggs’ and ‘can be trained’ appear to require specific contexts and the ability to verify the property may depend on the availability of a dynamic simulation for a surrounding context. In the ‘lays eggs’ property, the relationship to ‘owl’ was rated as less abstract (see Appendix C.) than the relationship to ‘termite’, perhaps because of the availability of concepts like ‘nest’ and ‘hatching’. Similarly, within ‘can be trained’, the relationship to ‘lion’ was rated as less abstract than ‘leopard’ perhaps because of simulations associated with ‘circus’ and ‘whip’. A neuroimaging design more sensitive to individual trials could differentiate the effect of contextual constraints on item-property relations a one source of abstractness. More abstract concepts, like ‘justice’ and ‘truth’, activate a frontotemporal network of cortical regions which may be driven by the necessity of generating and maintaining a sentential context for these items (Noppeney & Price, 2004).

As indicated in the neuropsychological literature, the prefrontal cortex would seem to be biased to represent knowledge that relies more on the language system than on perceptually-based properties. In support of this anterior-posterior dissociation, patients with damage to prefrontal areas are shown to have difficulty with verbal definitions. In contrast, individuals with category-specific impairments for identifying objects within a visual category, based on damage to ventral temporal regions, can still access semantic knowledge that relies more on verbal facts. As reviewed earlier, we have also found in related neuroimaging work that the verification of life status for plants and non-living natural kinds leads to increased activity in the prefrontal cortex. These types of instances initially cause confusions for young children and it is only through the acquisition of defining, and theory-based, features that they eventually learn the correct classifications. In this regard, the increased activation in the left inferior and middle prefrontal cortex may indicate the role of the language system in developing knowledge about the world. The left inferior prefrontal cortex would seem to support information about objects that cannot be directly experienced and so must necessarily coded through the language.

A curious result involves the increased activity in two of the observed prefrontal regions for the easier abstract verification question. When participants resolved whether an organism ‘can be trained’, these two prefrontal regions were associated with increased activity in contrast to the more difficulty verification of whether the items could ‘lay eggs’. This result suggests that the prefrontal cortex may fractionate into subnetworks specifically involved in some semantic functions but not others. While further work is necessary to test this prediction, the spatial relationship of these two regions being more lateral and anterior suggest another control function in processing abstract semantic knowledge, perhaps involving the binding of abstract properties to a referent. In contrast, the more medial and posterior regions show no such sensitivity to the difficulty-based differences and rather were associated specifically with the more abstract verification decisions. The prefrontal cortex may therefore subdivide into regions more involved in the representation of abstract semantic knowledge and regions more involved in the processing of that knowledge but assisted by the abstract representations or general processing mechanisms, like generating an extensional context, more likely to be engaged with abstract properties.

Based on these results, it appears that the prefrontal cortex may more generally support differing aspects of semantic processing depending on multiple factors. As also indicated by the results of Experiment 1, the prefrontal cortex would seem to operate as a second stage in semantic decisions depending on the specific content and processing requirements. The results of this second experiment indicate that the verification of abstract semantic properties specifically involves regions of the prefrontal cortex irrespective of behavioral difficulty. As suggested in the primate literature, regions of the prefrontal cortex would seem contribute the specific rule-based properties necessary for some semantic decisions but not others. In this regard, semantic memory would appear to consist of a highly distributed but also interactive network of brain regions with both local and global contributions to the representation and processing of world knowledge.

4. Summary and General Discussion

Across the two presented experiments, the results suggest three main conclusions regarding the cortical distribution and interaction of semantic knowledge. First, perceptual properties of objects appear to specifically rely on sensory brain regions used to encode those properties. That is, the retrieval of perceptual knowledge associated with an object appears to depend on the simulation of prior sensory experiences with the object. This result supports and extends perceptual symbol systems (Barsalou, 1999) and feature-based accounts for the cortical distribution of semantic memory (Martin & Chao, 2001) based on localized and specialized sensory regions in supporting perceptual properties of object knowledge. Domain-specific theories (Caramazza & Shelton, 1998), distributed representation accounts that assume no regional specializations (Tyler & Moss, 2001), and amodal representations (Caramazza et al., 1990) encounter greater difficulty accounting for these findings in explaining the organization of semantic knowledge.

Second, regions in the prefrontal cortex seem to specifically represent abstract properties of objects. These properties are likely coded in verbal terms and describe information about an object, in contrast to sensory-based representations of the object. This finding suggests a reconsideration of the role of the prefrontal cortex in semantic processing based on the abstractness of the task demands, either through a linguistic mediation of semantic processing (Noppeney & Price, 2004), especially when semantic associations become weaker (Wagner, 2001), or through more general cognitive control operations (e.g., Thompson-Schill, 1997) more necessary in semantic contexts. In this regard, the increased reliance on prefrontal mechanisms may represent a second stage in semantic processing as predicted by seminal models either through the specific content required (Smith, Shoben, Rips, 1974; Armstrong, Gleitman, & Gleitman, 1983) or more general decision processes (McCloskey & Glucksberg, 1979).

Finally, the controlled processing of semantic knowledge appears to rely on prefrontal regions depending on the role of those regions in representing modality-specific and abstract properties, or in directing top-down bias signals to posterior cortical regions (Miller & Cohen, 2001). Other prefrontal areas typically ascribed to more domain-general control functions (MacDonald et al., 2000, Schneider & Chein, 2003) would also seem to be involved in processing more difficult semantic decisions. This overall view suggests that semantic knowledge is distributed to specialized cortical regions based on the sensory modality involved and the abstractness of the semantic decision. In addition, semantic processing may be bound to more domain-general control operations with increasing demands of a given semantic decision. Prefrontal and sensory brain regions appear to interact to the extent that a given semantic decision becomes increasingly abstract and difficult, with specialized prefrontal regions supporting the different types of semantic processing based on the stimulus factors and decision context. How this distributed network interacts depending on multiple stimulus factors through specialized regions is the subject of much current debate regarding the role of the prefrontal cortex in cognitive control processes. Linguistic stimuli provide one way to characterize this executive control architecture while examining in detail how the prefrontal cortex, especially inferior regions, may fractionate based on the control mechanisms involved.

The role of the prefrontal cortex in the processing of linguistic information illustrates the numerous factors that likely influence cognitive control in this region. As with other domains, regions in the prefrontal cortex may represent content specific to language or more general processing mechanisms may be more likely utilized in service of linguistic functions. For instance, some of the first neuroimaging investigations of cognition (Petersen et al., 1988, 1989) found LIPFC activity when subjects generated a verb to a target noun in contrast to when nouns were simply repeated. Yet the difficulty in determining exactly how these two tasks differ illustrates the problem of identifying the factors of cognitive control that influence prefrontal mechanisms. As indicated in the current studies, the modality-specificity and abstractness of the stimulus materials are just two of many possibilities that require further testing.

The apparent material-specific delineation of anterior and posterior LIPFC regions arises from contrasts between phonological and semantic tasks (Bookheimer, 2002, Fiez, 1997, Gabrieli et al. 1998; though see Price et al., 1997), perhaps based on dedicated connections to posterior cortical areas (Gold & Buckner, 2002). Selective disruptions of these areas appear to specifically affect semantic and phonological decisions for word stimuli (Devlin et al., 2003). However, these results may more generally be explained in terms of the retrieval demands associated with each task (Barde & Thompson-Schill, 2002) rather than through domain-specific representations. Linguistic stimuli may be more likely to rely on these regions based on the specifics of naming items (Simons et al., 2001; Kan & Thomson-Schill, 2004) and in maintaining linguistic context (Kerns et al., 2004). Domain-general cognitive control mechanisms may be more likely to engage the LIPFC when linguistic tasks require sustained access to content in temporal regions (Noppeney, Phillips, & Price, 2004) and in suppressing irrelevant contextual information (Cardillo et al., 2004)

While many studies have examined the role of the LIPFC in the controlled processing of semantic knowledge, its exact role remains unclear. Numerous investigators have proposed that this region specifically supports semantic working memory (Petersen et al., 1989, 1988; Kapur et al., 1994; Demb et al., 1995; Martin & Chao, 2001) reflecting the predicted control function for the prefrontal cortex more generally with different domains localized to distinct regions. The effect of reducing memory demands by presenting repetitive semantic information supports this interpretation. Semantic priming effects lead to decreasing activity in inferior prefrontal areas (Bookheimer, 2002) with common regions for pictures and words (Wagner et al., 1997) but also a task-specific fractionization of this region (Wagner et al., 2000). Control mechanisms in the LIPFC may become increasingly necessary as semantic relations between lexical items become weaker. To test this prediction, Wagner and colleagues (Wagner et al., 2001) examined activity in LIPFC while the associative strength between a cue and target item was varied. The results indicate that the LIPFC was activated moreso by weaker associates suggesting that this region mediates a top-down bias signal that is recruited when semantic processing requires controlled processing, as predicted for the prefrontal cortex more generally (Miller & Cohen, 2001). However weaker semantic associations likely require controlled processing of stimulus independent representations, such as abstract properties, when the required decision is unclear and multiple representations competing representations must be mediated. Recent evidence (Smith et al., 1998; Patalano et al., 2001) suggests that prefrontal cortex may specifically represent learned decision rules used in the service of task demands and in controlling prepotent tendencies, as the response properties of neurons in the primate analogue indicate (Wallis et al., 2001). Category and property decisions that require access to such abstract semantic information specifically activate the ventrolateral prefrontal cortex (Noppeney & Price, 2004) and the presented results similarly support this interpretation.

In addition to suggesting that specific types of properties are processed through prefrontal mechanisms, these regions may support more general control operations that dissociate based on input from posterior areas (Wilson et al., 1993; Romanski & Goldman-Rakic, 2002). That is, rather than mediate semantic content per se, the prefrontal cortex may be more specifically involved in particular control functions such as the selection of task-relevant responses from among competing alternatives (Thompson-Schill et al. 1997, 1999). For instance, Thompson-Shill and colleagues (1997) varied semantic selection demands by manipulating the number of distracters in classification, comparison, as well as verb generation tasks. Items with many possible responses, in contrast with fewer possible alternatives, led to increased activity in the LIPFC irrespective of the task demands or stimulus type. Activations generally found in ventrolateral prefrontal cortex during both classification and generation tasks further suggest the role of multiple generic control processes (Cabeza & Nyberg, 2000) instead of specific types of representations. The present results support both material-specific and domain-general interpretations for the role of the prefrontal cortex in the representation and processing of semantic knowledge. Future studies, by controlling for the modality and abstractness of task demands, are likely to distinguish between these possibilities.

The results of the presented experiments support a feature-based reduction of semantic memory in which the knowledge of specific attributes of objects depends on brain regions specialized to represent and process those properties. In this regard, knowledge representation can be reduced to the constituent attributes of objects and the ways that particular aspects of objects are encoded. During perceptual experiences of objects, particular sensory brain mechanisms are used in processing the experience. For instance, in hearing a dog’s bark or a cat’s meow, sensory mechanisms in primary and secondary auditory must become active. The results indicate that recalling knowledge of these sounds involves increased activation of auditory brain regions. Similarly, visual, tactile, and gustatory knowledge of objects appear to rely on cortical regions specifically involved in encoding each of these sensory modalities. This direct perceptual knowledge of an object would therefore seem tied to the varied sensory experiences that gave rise to it, perhaps through increased specificity of object simulations with increased exposure to particular objects (Barsalou, 1999).

In contrast, knowledge of objects can also be encoded more indirectly through the language system especially when multimodal, and summative, representations are required. A teacher, parent, friend, or even a book can describe objects in ways that induce the sensory representations of the object. However this verbal format can also provide new knowledge about an object that describes abstract properties and situational contexts that cannot be directly experienced. Scientific knowledge may therefore develop *about* objects in a way that describes properties that become associated with the representation *of* the object. For instance, the example of how biological knowledge develops again provides a useful focusing point. Young children are initially confused by the overt behavior of plants and non-living things such as a cloud, a truck, and a rock, concluding that the plants are not alive while the latter examples are alive. The difficulty for the child seems to involve learning to ignore the superficial perceptual features of objects in favor of underlying, and unseen, biological mechanisms and their associated, rule-based, properties like those involving reproduction, growth, and eventually concepts of cellular division and genetics. Current feature-based models do not presently address the role of such features, or the cognitive control necessary to mediate them, in accessing semantic knowledge. The results suggest that the prefrontal cortex is specialized to the represent and process such abstract properties, and perhaps through a top-down modulation of processing of multiple modality-specific properties in posterior brain regions. Given the role of the prefrontal cortex in memory encoding and its delayed maturation relative to sensory regions (Diamond & Goldman-

Rakic, 1991), these anterior association areas would seem ideally suited to the representation and processing of abstract semantic knowledge as language develops.

Based on seminal and contemporary categorization models, it should be no surprise that the brain takes advantage of feature similarity to ground the structure of particular categories. More difficult are borderline cases when it is not clear, from perceptual features alone, whether an item is a member of the given category. Consider for instance, whether an ‘apple’ (or still more difficult, your ‘ear’) is a living thing. If we were to simply group objects based on how they look, or even their near associates in semantic space (e.g., ‘pie’ or ‘sauce’), learning scientific concepts (i.e., ‘cell division’) would be trivial. However the conflict of perceptual features prevents easy categorization. For example, the neuropsychological and neuroimaging literatures have tended to treat the living/nonliving distinction as a rather obvious division of the world, even as one based on innate domain-based mechanisms. But for the young child this distinction is anything but simple.

In a domain as perceptually diverse as biology (consider: ‘cow’, ‘daisy’, ‘wasp’, and ‘amoeba’), children must begin to infer abstract causal mechanisms, and often without a firm perceptual grounding, to explain superficial similarities that are initially misleading, as in linking ‘plants’ and ‘animals’ under the common category ‘living thing’ (Carey, 1985; Keil, 1989). The category of plants is so perceptually dissimilar from that of animals that young children initially claim that only animals are living things. The young child latches onto self-generated movement as a reliable cue as to whether something is alive (for a recent review see Gelman & Opfer, 2002). Abstracting from visual input, and perhaps using simple functional features, would seem to require the acquisition and coordination of multimodal rules (e.g., all living things ‘die’ and ‘reproduce’) with these perceptual cues. The young child will eventually need to explain outward movement in terms of an underlying biological mechanisms, and with a few years of experience and education they'll attach concepts like 'growth', 'nutrition', and eventually 'cell division' to this skeletal structure. Likely, if you're still struggling with whether an 'apple' is alive, you've seen that the similarity of perceptual features alone is not sufficient in drawing category boundaries, and thus you've been busy retrieving facts and concepts from semantic memory in a concerted effort to find the answer. Such information is crucial when deciding category membership for more borderline cases. The interaction between posterior and prefrontal brain regions may provide an instantiation of seminal two stage models of category structure and processing. The boundaries of a given category, especially with regard to borderline cases, are likely processed using areas in prefrontal cortex to control semantic retrieval and selection likely through the representation in these same areas of abstract properties. In this way, seminal two-stage models would seem to be instantiated with posterior regions of cortex performing the general feature pattern match of the first stage, while more defining, and relational features are processed on-line as necessary to resolve more borderline decisions.

Neuroimaging results of category-specific effects provide further support for the dissociation of semantic categories to particular regions of cortex. The form-function correlations that give rise to these effects are grounded in a highly distributed and overlapping cortical network reflecting vectors in, more specifically, a semantic feature space. The maximal activation of just one type of feature to a particular region is likely insufficient to capture the myriad of ways that semantic memory is coded. Recent neuroimaging work (Chao, Weisberg, & Martin, 2002) further supports the view that the content of category-related activations are based on specific sets of semantic features and mediated by the functional architecture of their near neighbors (i.e., the neural basis of conceptual similarity). Such results suggest that this distributed basis for knowledge representation is the direction of future neuroimaging investigations into semantic memory. For instance, the representation of an item and information about an item can differ significantly based on the surrounding context and goals (Barsalou, 1991). You might take a ‘blanket’ to the beach as something to rest on, but its constituent functional use (i.e., to keep you warm) is the generally intended usage. In this regard, semantic knowledge is inherently flexible, depending on contextual circumstances and constraints. The processing of semantic features and their relations and extensions become crucially important in any attempt to understand the breadth and depth of semantic knowledge.

In examining general theories of category structure and processing, the brain appears to make no clear distinction between perception and conception, as predicted by more recent categorization models (Goldstone & Barsalou, 1994). That is, semantic features would seem to be located in or near the sensory-motor areas of cortex in which the perceptual experiences were initially encoded, as predicted by the perceptual/functional hypothesis. Semantic categories, based on converging evidence from neuropsychology and neuroimaging, likely emerge from the feature intercorrelations that occur in processing similar items through shared cortical regions (Simmons & Barsaou, 2003). In addition, the breakdown patterns of some patients seem to further support this organization as prototypes remain somewhat intact (e.g., De Renzi & Lucchelli, 1994), even though distinguishing, and more peripheral features, are lost which would normally serve to differentiate near neighbors in semantic space.

As a whole these neuroimaging results strongly suggest that, general semantic processing and, more specifically, conceptual knowledge of categories, such as animals and tools, is instantiated by distinct anatomical regions. In this regard, the neuroimaging of semantic memory provides support for the dissociations of category-specific deficits found in the neuropsychological literature. In addition these findings provide further support for recent proposals of perceptually-based symbols and the distinction between the representation of an object and information about it. This dissociation between perceptual features used in the direct experience of an object and abstract relational properties, coded in verbal terms, used to indirectly characterize an object, provides a further extension to current feature-based models. Categories of objects, and effects specific to them, can therefore arise from perceptual similarity and/or more abstract defining features based on the cortical distribution and interaction of semantic knowledge.

TABLE HEADINGS

Table 1: Design of Experiment 1 with questions and example affirmative stimuli

Table 2: Neuroimaging results of Experiment 1 based on conjunction analyses for each semantic modality and difficulty-based regions at $p < .01$ (uncorrected). Listed are the approximate anatomical locations, Brodmann Areas (BA), and Talairach coordinates, as well as the number of voxels in each ROI and its average statistical value.

Table 3: Design of Experiment 2 with verification questions.

Table 4: Neuroimaging results of Experiment 2 based on a direct contrast between the semantic verification blocks and the nonword letter detection task at $p < .05$ (Bonferroni correction). Listed are the approximate anatomical locations, Brodmann Areas (BA), and Talairach coordinates, as well as the number of voxels in each ROI and its average statistical value.

Table 1: Design of Experiment 1

Difficulty	Easy	emerald	thunder	cotton	caramel
	Hard	cabbage	grenade	flannel	coconut
		GREEN?	LOUD?	SOFT?	SWEET?
		Visual	Auditory	Tactile	Olfactory/Gustatory
		Sensory Modality			

Questions:

Does X look green?

Does X sound loud?

Does X feel soft?

Does X taste sweet?

Table 2: Neuroimaging results of Experiment 1

Analysis	Region	BA	Talairach Coordinates			Voxel count	p value
			x	y	z		
Visual	L Superior Parietal	7	-35	-71	45	125	0.001708
	L Middle Temporal	37	-52	-58	0	10	0.001662
	L Middle Temporal	37	-59	-46	-6	10	0.003038
Auditory	L Superior Temporal Sulcus	22 / 42	-56	-48	7	39	0.002674
Tactile	L Post-central Gyrus	2 / 40	-54	-33	34	83	0.002254
	L Pre-central Gyrus	4 / 6	-47	5	25	142	0.002449
	L Premotor	6 / 9	-34	29	20	165	0.002458
Gustatory	L Inferior Frontal	11 / 12	-19	29	-7	5	0.002997
Hard	L Dorsolateral Prefrontal	9	-50	11	31	75	0.004550
	Anterior Cingulate	32	-1	8	50	57	0.004098
	L Inferior Prefrontal	46	-41	32	11	42	0.004148

Table 3: Design of Experiment 2

Difficulty	Easy	Does X have four feet?	Can X be trained?
	Hard	Does X have fur?	Does X lay eggs?
		Sensory	Abstract
		Abstractness	

Table 4: Neuroimaging results of Experiment 2

Analysis	Region	BA	Talairach Coordinates			Voxel count	p value
			x	y	z		
Words - Nonwords	L Inferior Prefrontal Cortex	47	-31	35	0	25	5.44E-08
	L Inferior Prefrontal Cortex	47	-48	37	2	11	4.15E-07
	L Middle Frontal Gyrus	9	-44	21	32	56	2.14E-08
	L Middle Frontal Gyrus	10	-33	50	12	66	2.73E-08

FIGURE HEADINGS

Figure 1: Average correct response times for each of the property verifications across the four sensory modalities and the letter detection control.

Figure 2. Average correct response times for each of the property verifications across the four sensory modalities and divided by easy and hard trials defined a priori.

Figure 3. Average correct response times for each of the property verifications across the four sensory modalities and segregated by easy and hard trials based on data from the neuroimaging participants.

Figure 4. The visual property verification specifically activates the left ventral temporal cortex.

Figure 5. The auditory property verification specifically activates the left superior temporal sulcus.

Figure 6. The tactile property verification specifically activates regions in the left somatosensory, primary motor, and pre-motor cortices.

Figure 7. The gustatory property verification specifically activates the left orbitofrontal cortex.

Figure 8. The tactile and gustatory verifications specifically activate the left inferior prefrontal cortex irrespective of item difficulty.

Figure 9. The left dorsolateral prefrontal cortex shows an interaction of the probed semantic modality and the difficulty of the presented items.

Figure 10. The anterior cingulate cortex shows a main effect of response difficulty.

Figure 11. Accuracy for verification questions based on rated abstractness and difficulty.

Figure 12. Average correct response times to each of the verification questions based on rated abstractness and the average difficulty.

Figure 13. The abstract verifications led to increased activity in a medial aspect of BA 47 of the LIPFC irrespective of response time.

Figure 14. A lateral region in BA 47 of the LIPFC was responsive to both question abstractness and difficulty, though driven by the easier abstract verification.

Figure 15. An anterior aspect of the left middle frontal gyrus responded most to the easier abstract verification suggesting a subnetwork based on some other semantic function.

Figure 16. The left DLPFC was especially sensitive to item-property abstractness with no differences based on behavioral performance.

Figure 1.

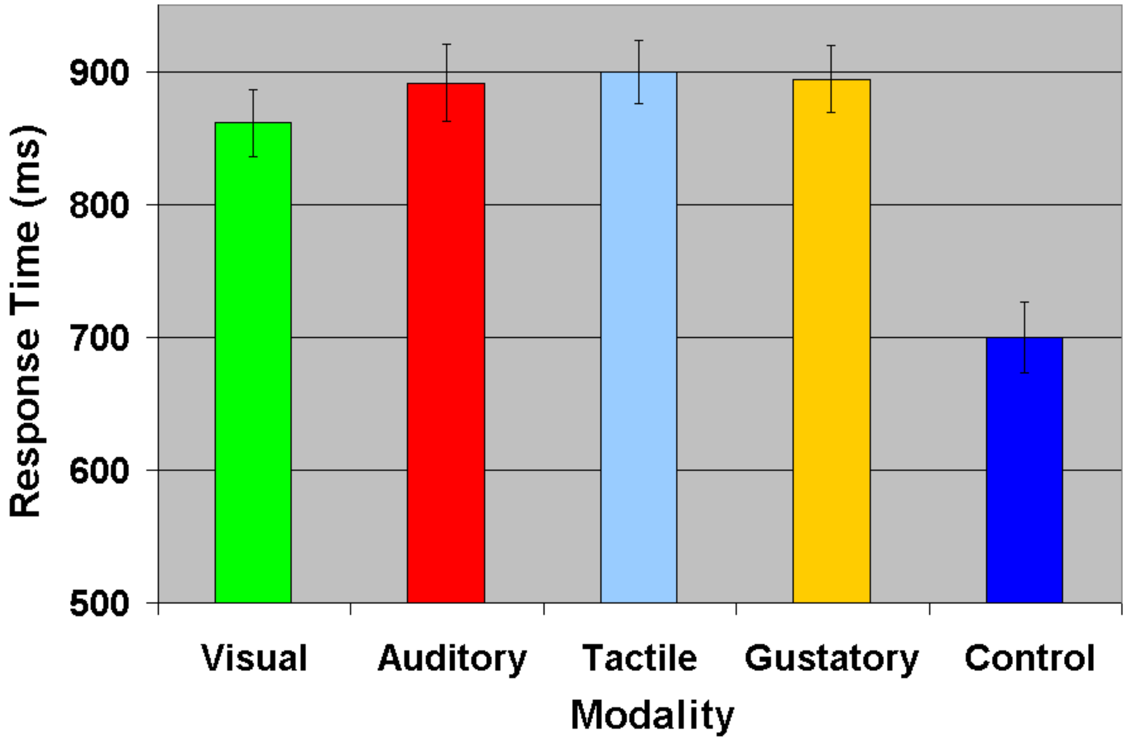
Response Time (ms)
900
800
700
600
500
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 2.

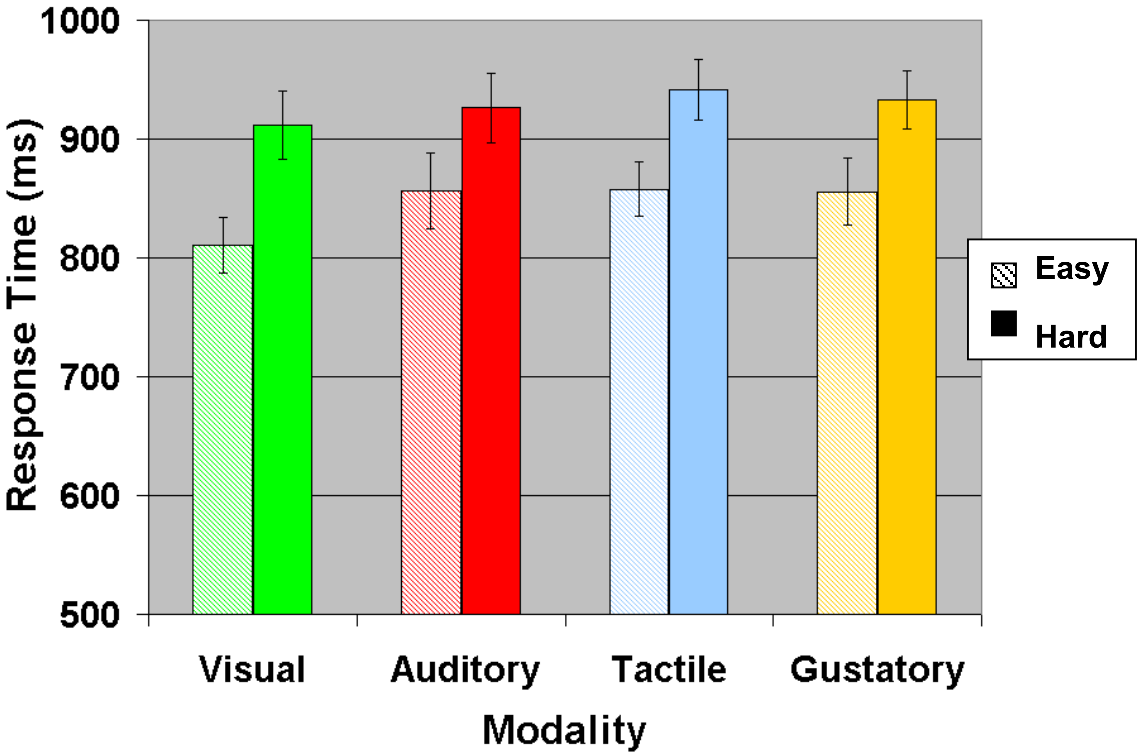
1000
900
800
700
600
500
Response Time (ms)
Visual
Auditory
Tactile
Gustatory
Modality
Easy
Hard

Figure 3.

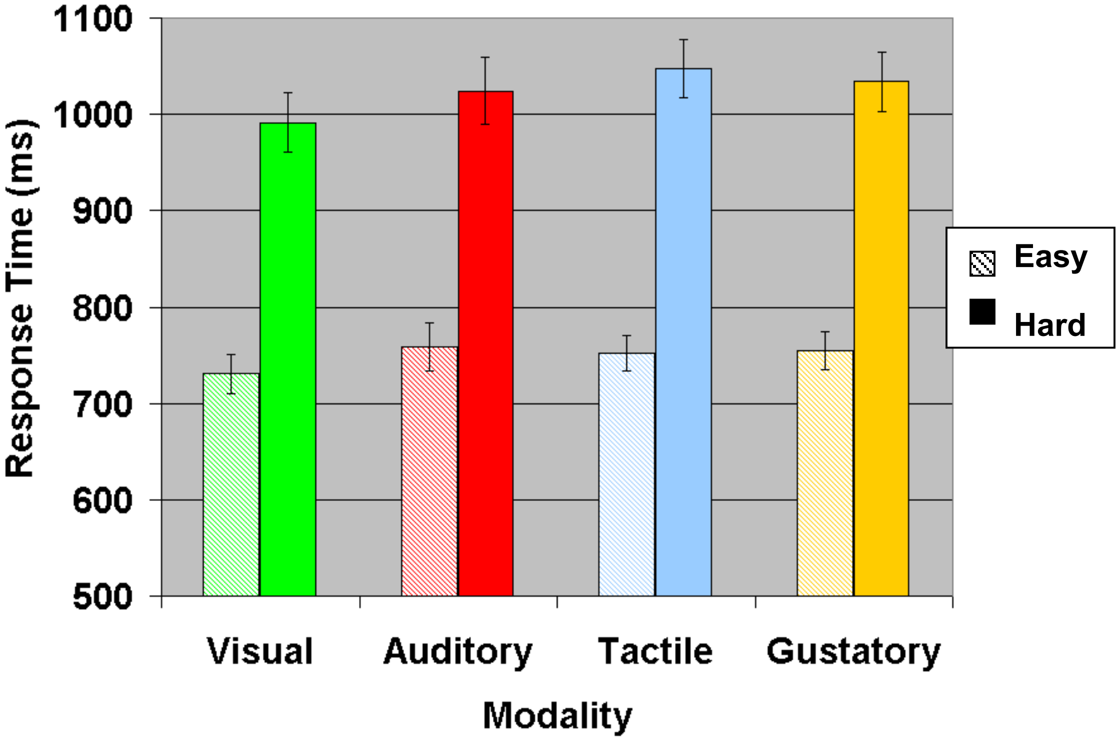
1100
1000
900
800
700
600
500
Response Time (ms)
Visual
Auditory
Tactile
Gustatory
Modality
▨ Easy
■ Hard

Figure 4.

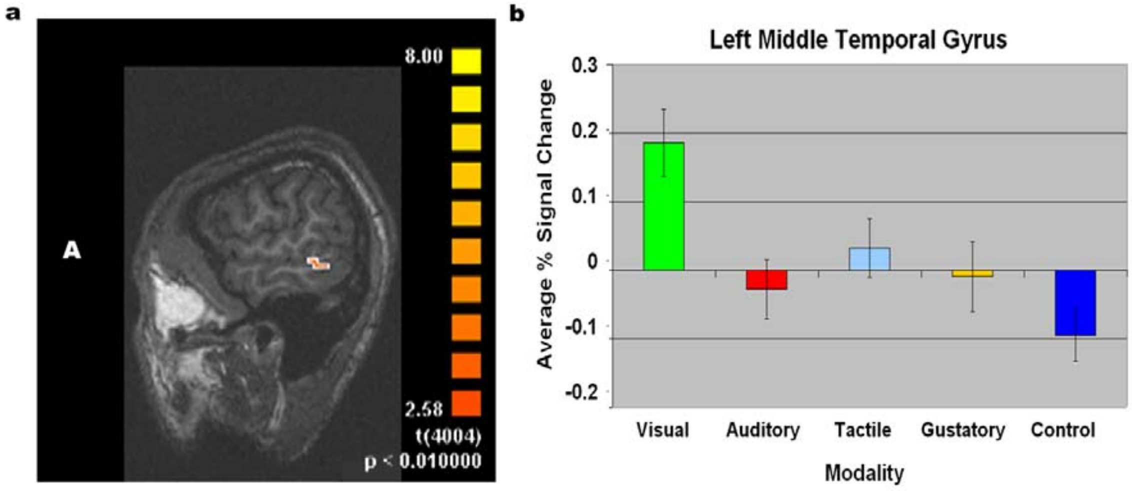
a
A
8.00
2.58
t(4004)
p < 0.010000
b
Left Middle Temporal Gyrus
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 5.

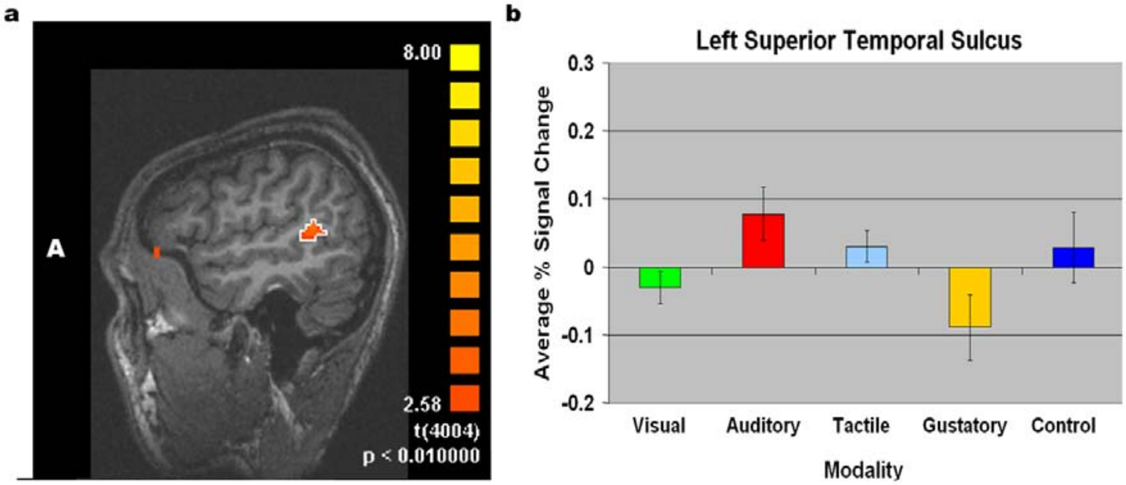
a
A
8.00
2.58
t(4004)
p < 0.010000
b
Left Superior Temporal Sulcus
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 6.

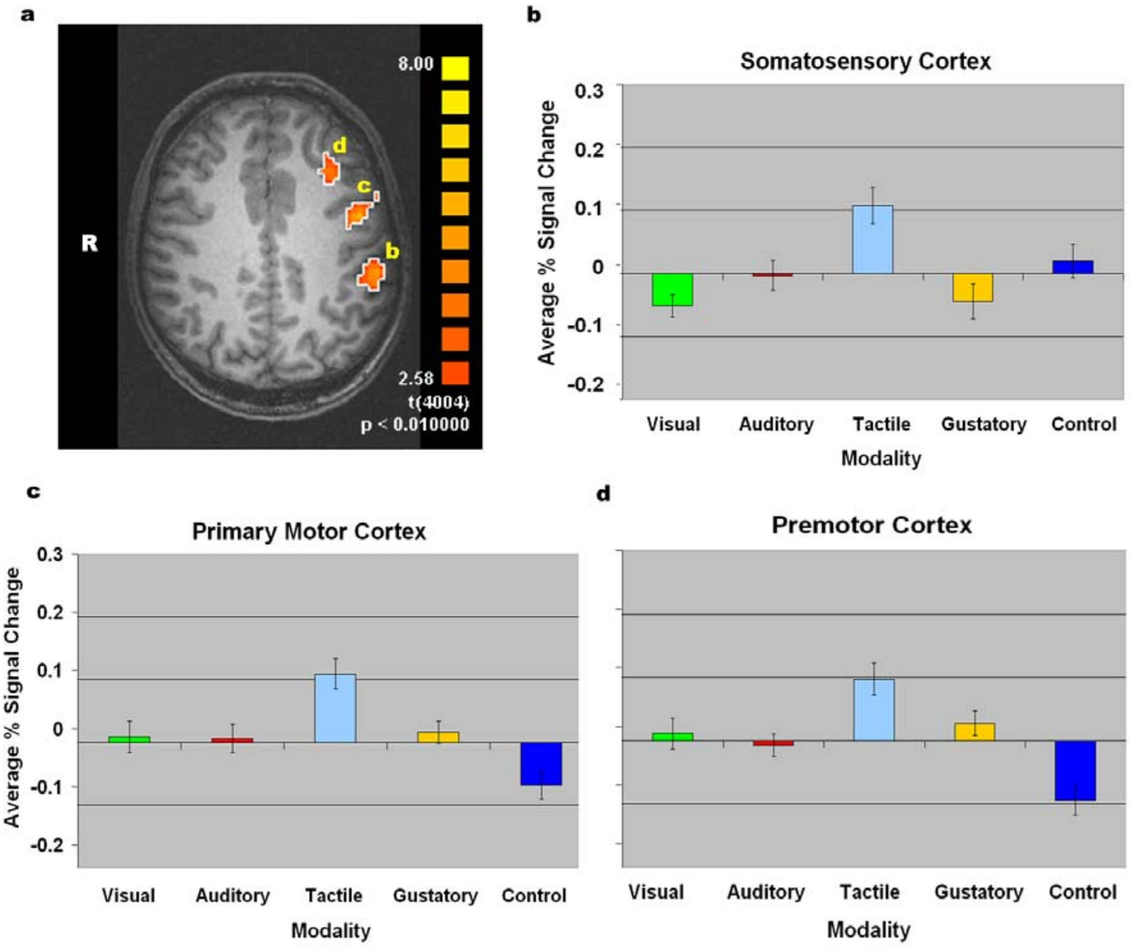
a
R
8.00
d
c
b
2.58
t(4004)
p < 0.010000
b
Somatosensory Cortex
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Visual
Auditory
Tactile
Gustatory
Control
Modality
c
Primary Motor Cortex
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Visual
Auditory
Tactile
Gustatory
Control
Modality
d
Premotor Cortex
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 7.

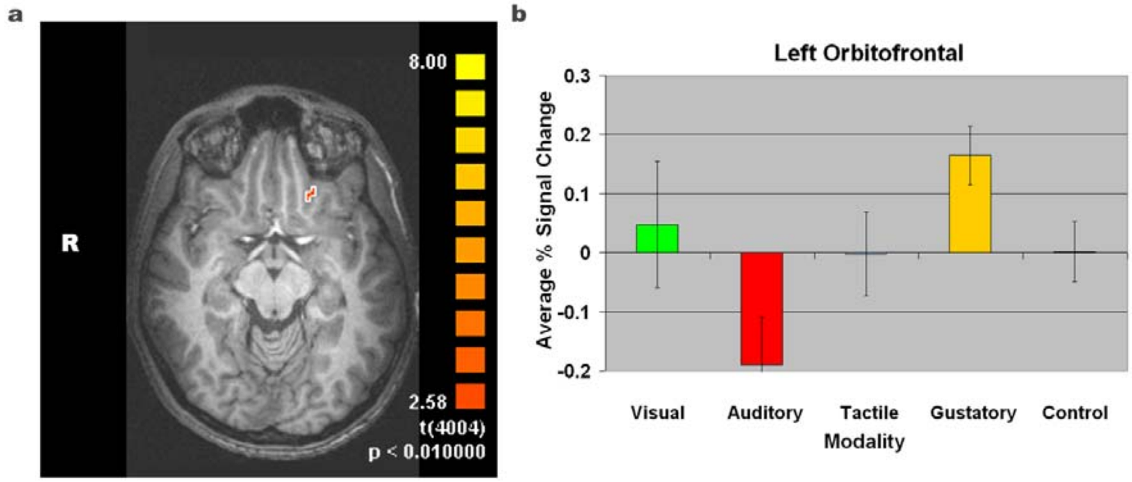
a
R
8.00
2.58
t(4004)
p < 0.010000
b
Left Orbitofrontal
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 8.

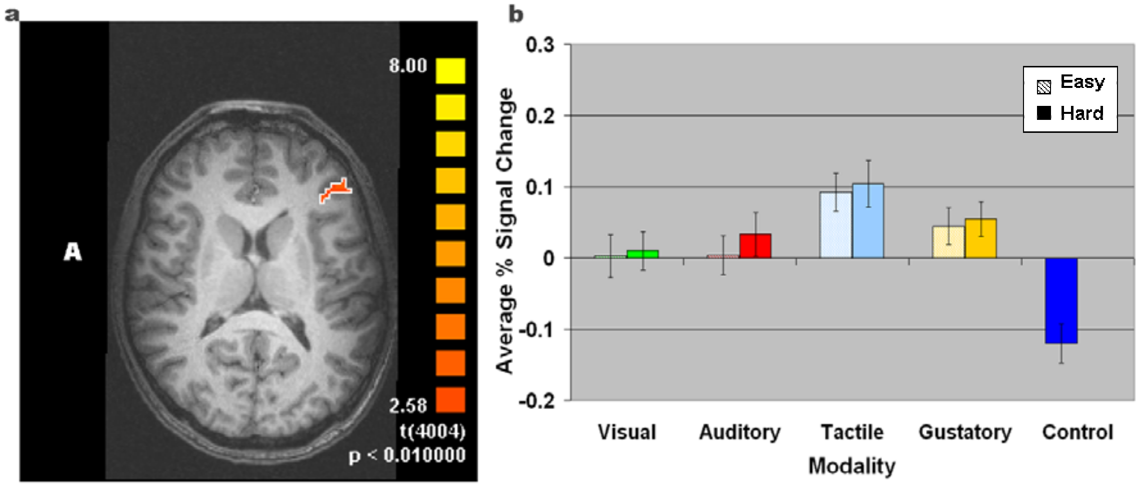
a
A
8.00
2.58
t(4004)
p < 0.010000
b
0.3
0.2
0.1
0
-0.1
-0.2
Average % Signal Change
Easy
Hard
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 9.

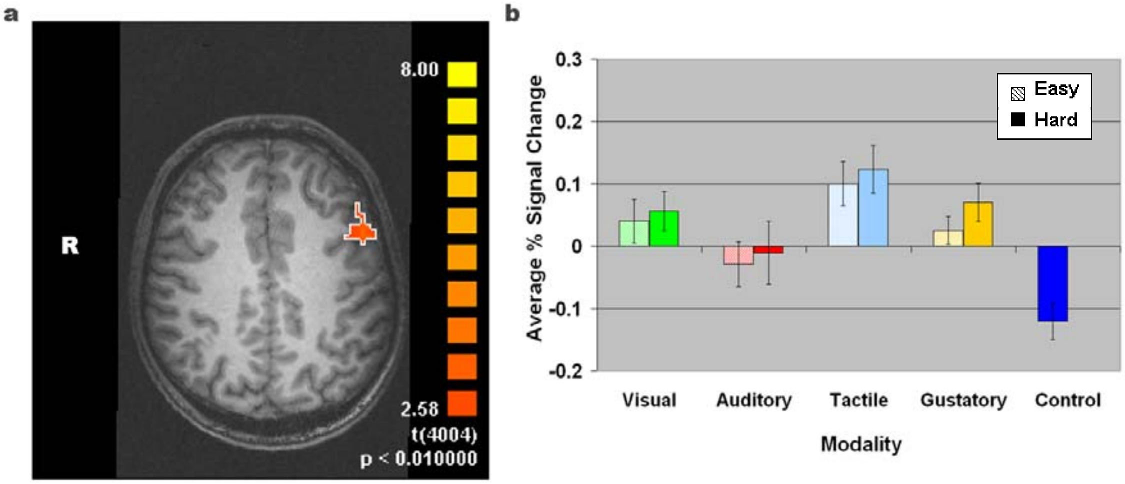
a
R
8.00
2.58
t(4004)
p < 0.010000
b
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Easy
Hard
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 10.

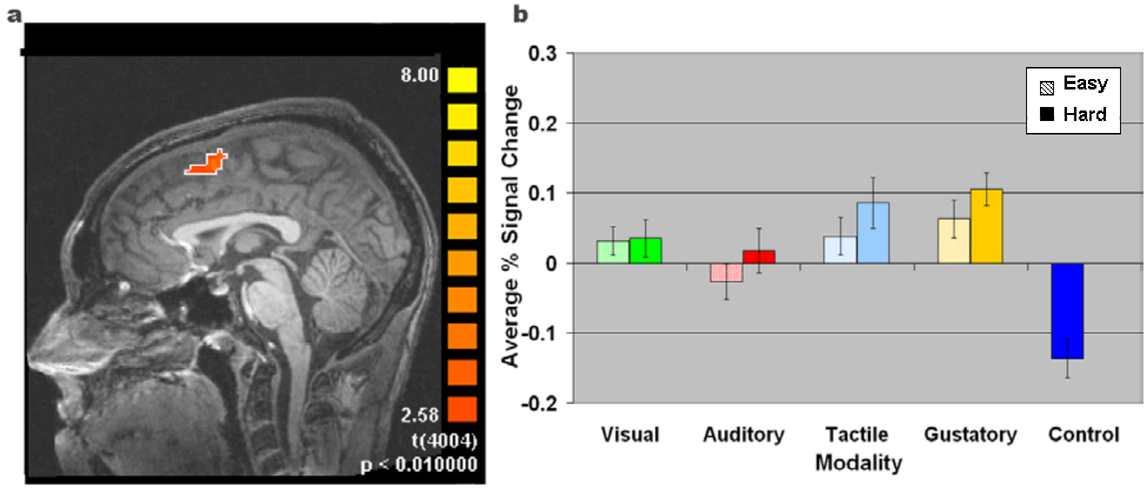
a
8.00
2.58
t(4004)
p < 0.010000
b
0.3
0.2
0.1
0
-0.1
-0.2
Average % Signal Change
▧ Easy
■ Hard
Visual
Auditory
Tactile
Gustatory
Control
Modality

Figure 11.

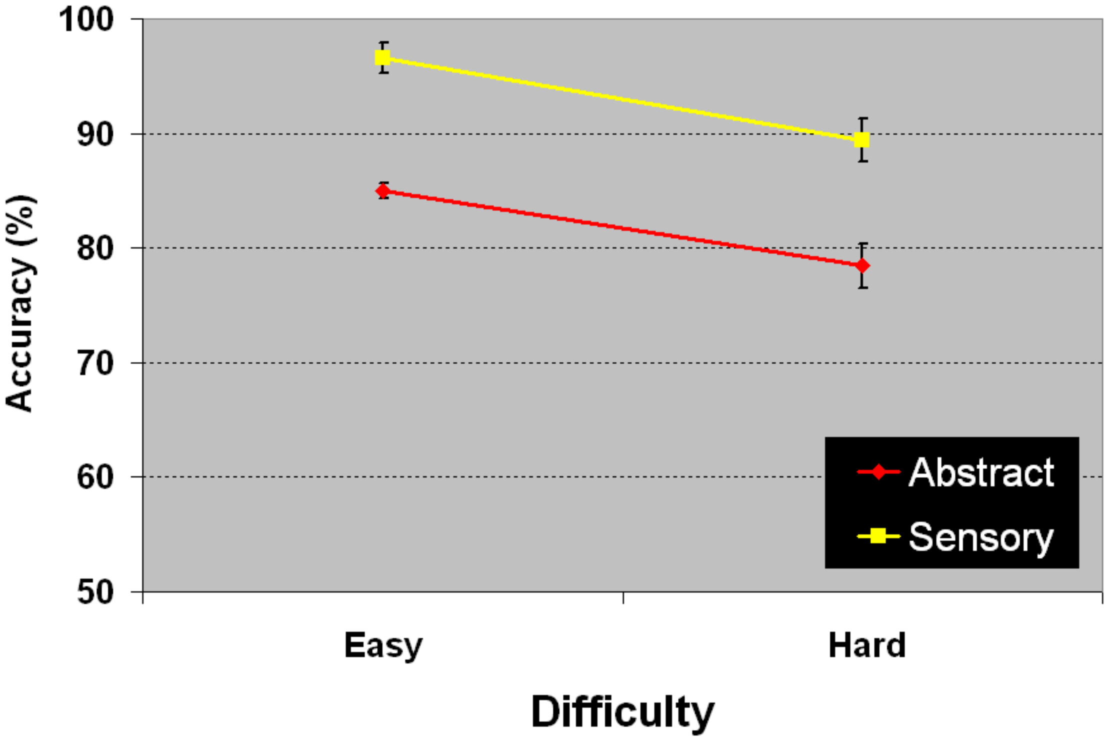
100
90
80
70
60
50
Accuracy (%)
Abstract
Sensory
Easy
Hard
Difficulty

Figure 12.

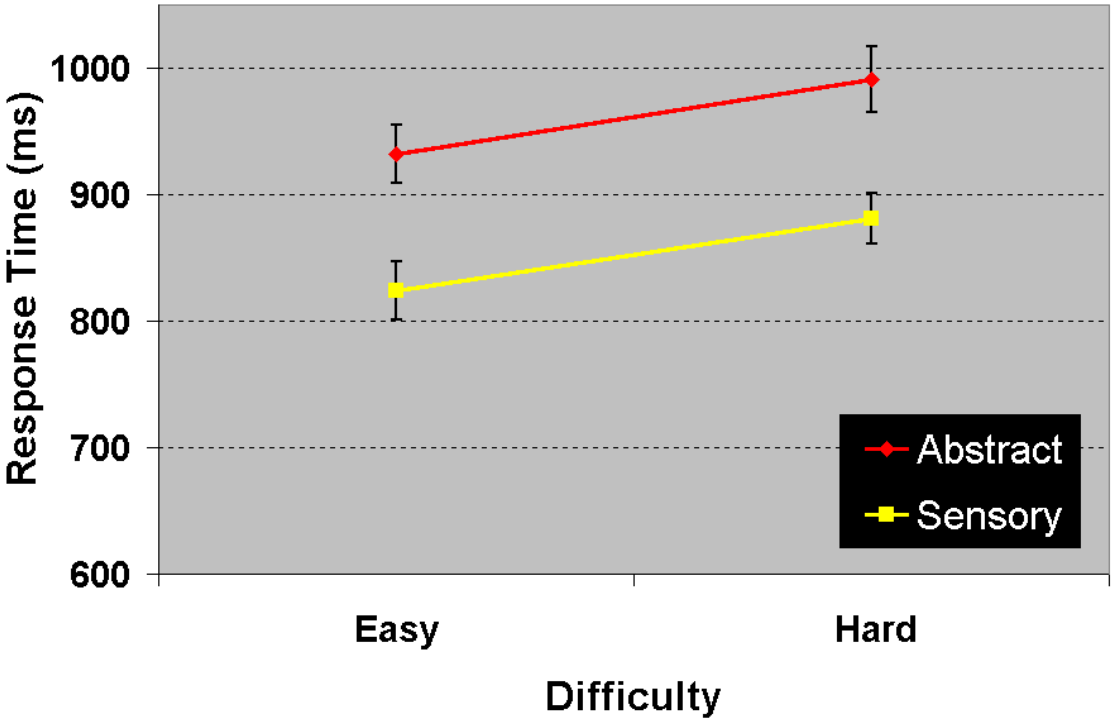

Response Time (ms)
1000
900
800
700
600
Easy
Hard
Difficulty
Abstract
Sensory

Figure 13.

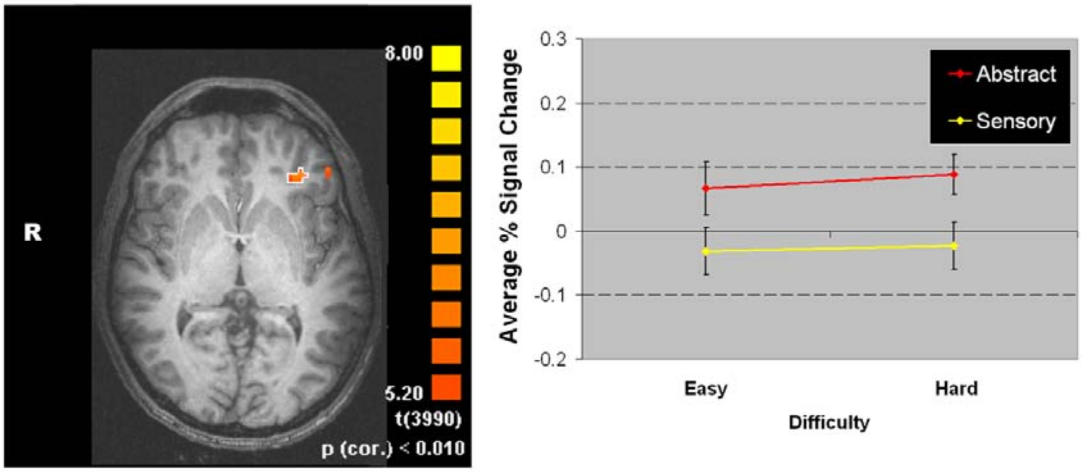

R
8.00
5.20
t(3990)
p (cor.) < 0.010
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Abstract
Sensory
Easy
Hard
Difficulty

Figure 14.

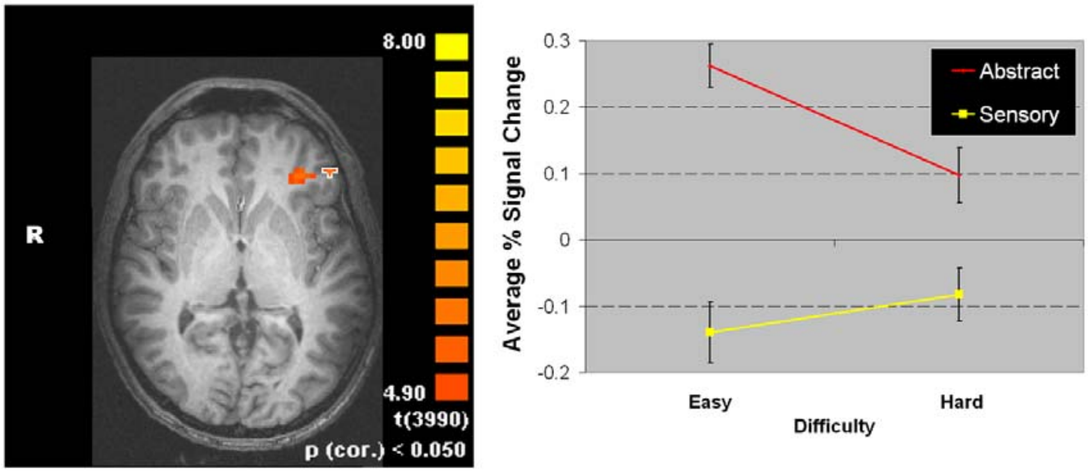
R
8.00
4.90
t(3990)
p (cor.) < 0.050
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Abstract
Sensory
Easy
Hard
Difficulty

Figure 15.

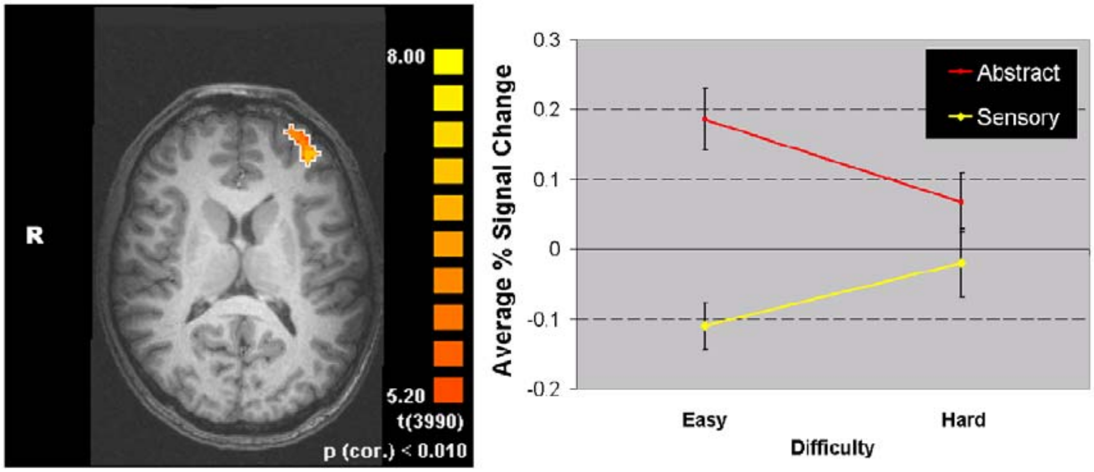
R
8.00
5.20
t(3990)
p (cor.) < 0.010
Average % Signal Change
0.3
0.2
0.1
0
-0.1
-0.2
Abstract
Sensory
Easy
Hard
Difficulty

Figure 16.

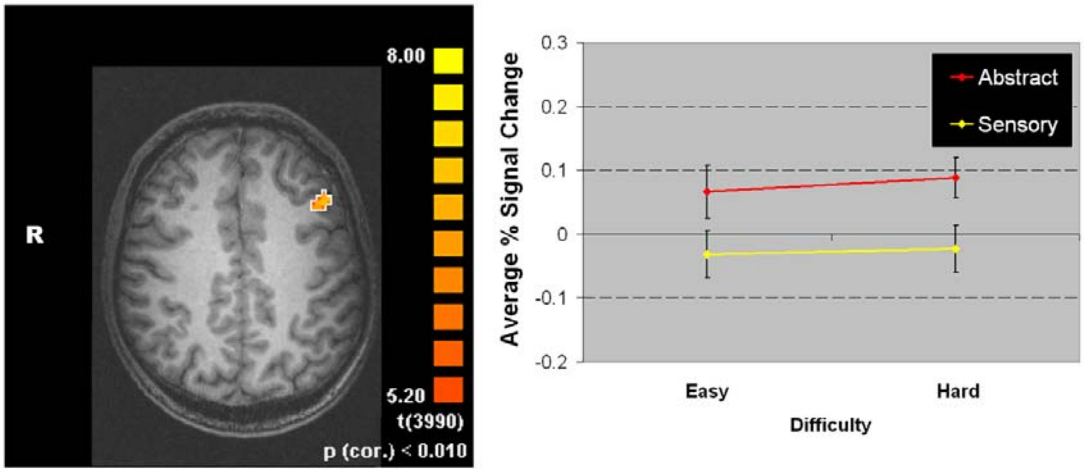
R
8.00
5.20
t(3990)
p (cor.) < 0.010
0.3
0.2
0.1
0
-0.1
-0.2
Average % Signal Change
Abstract
Sensory
Easy
Hard
Difficulty

APPENDIX A - Neuroimaging Methods

The underlying temporal resolution of fMRI is limited by the hemodynamic response of the Blood Oxygen Level Dependent (BOLD) technique (Kwong et al., 1992; Bandettini et al., 1992). BOLD is sensitive to changes in blood flow and oxygenation that typically develop in two seconds and peak in 5-7 seconds. The activation signals appear to be nearly linear and additive (Buckner et al., 1996; Dale & Buckner, 1997). The fMRI analysis is inherently a difference technique that requires a control and contrasting conditions. The MR signal is very heterogeneous across the cortex, with large variations in signal (e.g., between white matter and cerebral spinal fluid). The small changes (e.g., 1%) in signal due to brain activation are typically revealed only after subtracting activation found in the visuo-motor control condition, though we expected main effects between semantic conditions with less than 1% signal change. The visuo-motor control task was selected to enable us to factor out effects due to simply seeing a stimulus, selecting a decision, and making the appropriate motor response.

The true benefit in using event-related designs is the ability to align the recorded behavioral and fMRI data on any stimulus or response event. Events types (e.g., experimental and control, stimulus response, stimulus accuracy) can then be combined or subtracted. Statistical tests specify a temporal time window for assessing the significance of activation using GLM-based contrasts either between conditions or between the differences of each condition relative to a baseline (e.g., local baseline per stimulus period). This produces a significance difference test for each voxel at typically 250,000 cortical points.

APPENDIX B – Behavioral Norming Experiments

For both of the presented experiments, norming experiments with behavioral-only subjects were used to: 1) identify verification questions that vary according to ‘abstractness’; 2) identify verification questions that vary according to ‘difficulty’ and to identify clear affirmative and negative items within each question; and 3) match each question for possibly confounding lexical factors. Each of these factors was computed in several groups (n = 23 - 30) of behavioral participants tested at separate times under speeded conditions at personal computer terminals. The procedures, as described below, within each norming experiment were held constant whereas norming tasks across stimulus dimensions were varied to control for the factor of interest.

1) Rating item-property abstractness

In order to be sure that Experiment 2 controlled for the differences between abstract and perceptual properties, and so that there was no abstractness difference between verification questions of Experiment 1, for all of the verification decisions presented in Experiments 1 and 2, behavioral participants were asked to rate the abstractness of item-property relations. Items were presented in sequences while the property was held constant with items and properties varying randomly. Participants were asked to rate, on a 5-point scale, whether they knew the item-property relationship through sensory or abstract facts. The following instructions were read to every participant:

> In this experiment you will be asked to decide about how you know certain facts. Some facts can be known by having directly experienced them, whereas other facts are learned indirectly either through reading about them or being taught them. **Sensory** facts are facts that you know from seeing, hearing, smelling, tasting, or touching things. For instance, the fact that most animals have eyes is obviously based on what the animal looks like. This is a sensory fact – that it can be readily learned from sensory experience. **Abstract** facts, on the other hand, are learned by either reading about them or by having been taught them. For instance, you

> know that all living things have cells. This is an abstract fact – that you must learn it indirectly either by reading about it or by being told it by someone else.
>
> Now, sensory and abstract facts may refer not only to things like 'eyes' and 'liver', but also to the relationship between a thing and its properties. For instance, whether a 'tiger' 'has eyes' is a sensory relationship. However, whether a 'worm' 'has eyes' is likely an abstract relationship, based on how you learned these facts. That is, it is not clear from how worms look whether they have eyes – you would have to learn this fact either through a book or another person and so it is an abstract fact.
>
> This experiment will ask you to judge the relationship between things and a given property, for example 'tiger' and 'has a liver'. Your job will be to determine how you do, or would, know, whether an item has the given property based on a 5-point scale. If you know the item has the property from sensory experiences with the item, then this is a sensory relationship and you'll press the #1 button. If you know the item has the property from being taught the fact, then this is an abstract relationship and you'll the #5 button. If you think your knowledge based somewhere exactly in between sensory and abstract facts, press the #3. However, try to go to one side with the #2 representing mostly sensory facts and the #4 representing mostly abstract facts.
>
> Keep in mind that you aren't trying to answer whether the item has the property or not. You're only indicating what type of facts, sensory or abstract, you would use to answer that question. So it's not important whether a worm actually has eyes or not, only what types of facts you would use to arrive at the answer.

Across three versions of this experiment, 15 properties, five in each version of the experiment, were tested where only they and the probed items varied. Some items were repeated multiple times across the set of tested properties but none were tested more than once for each property. Participants had 7 seconds to respond and were presented with between 700 and 800 trials in each version of the experiment. Any participant that responded in less than 500 ms for over 10% of trials was removed from further analysis. Based on this criterion, the data of 5 or 6 participants was not included in the final behavioral norms, which were based on data from 17 to 20 participants.

The results of this experiment produced verification questions that ranged on this abstractness scale, across all presented items, from 1.75 (tastes sweet) to 4.10 (has a liver). By including and excluding items from the set presented in the respective neuroimaging experiments, it was possible to match the verification questions on this dimension for Experiment 1 and to vary the questions on this factor for Experiment 2.

2) Verifying properties to items

To be sure that no items were particularly difficult within each verification question of Experiments 1 and 2, all questions were piloted with a superset of items used in the neuroimaging experiments. This testing of the stimulus materials replicated the overall procedure and design of the neuroimaging experiments with behavioral-only participants told to verify whether each presented items possessed the given property. Across participants, average response times and the average response, affirmative or negative, for each item was computed within each verification decision. These scores were then used eliminate items judged inconsistently and to match affirmative and negative items within each property as well as to match for these factors across the set of questions. In the regard, this data was used to generate easy and hard items within the modality-based property verifications of Experiment 1 and to control easy and hard verification questions for Experiment 2.

3) Attributing lexical status to items

The lexical familiarity measure was calculated from response times and accuracy patterns of behavioral-only participants when the items are presented in a lexical decision experiment. This effort has aimed to derive a direct empirical measure for how well word items are known, in contrast to text-driven frequency estimates (Francis & Kucera, 1982), which are susceptible to corpora selection bias and inconsistencies between how words are used in speech and in print. Lexical decision data for about 30 subjects for over 600 words have been transformed into a standardized z-space that represents a word familiarity index. This index reflects the ease with which participants ascribe lexical status (e.g., ‘algae’ = -1.15 whereas ‘apple’ = +1.05 – with the most positive being responding to very quickly and very accurately) and therefore empirically controls for how well a given word stimulus is known. Our previous neuroimaging efforts have used this measure to more strongly rule out lexical factors, like familiarity or letter length, in activation differences.

APPENDIX C – Items used in Experiments 1 and 2

Key.

Affirmative – Items to which a ‘yes’ response was required

Negative – Items to which a ‘no’ response was required

Stimulus – Items used in the block

Familiarity – Lexical familiarity score derived from response times and accuracy for the item when presented in a lexical decision experiment.

Letters – Number of letters in the item

Syllables – Number of syllables in, or phonological length of, the item

Response – Average response given, where yes = 1 and no = 2, when item was verified for the property by the behavioral-only participants.

RT - Average response time for the item when the property is verified by the behavioral-only participants.

Abstract - Average abstractness rating for the item-property relationship, from 1 to 5 where 1 is based on sensory facts and 5 is based on abstract facts.

Experiment 1 items. Visual modality (*Looks Green?*)

Green		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Easy Affirmative	1	basil	-0.12	5	2	1.20	696.80	1.50
	2	celery	0.26	6	3	1.00	706.68	1.15
	3	crocodile	0.27	9	3	1.04	686.24	1.30
	4	cucumber	0.34	8	3	1.08	647.38	1.20
	5	emerald	0.04	7	3	1.04	687.67	1.50
	6	lettuce	0.45	7	2	1.00	661.48	1.20
	7	lizard	0.69	6	2	1.04	699.24	1.20
	8	mint	0.21	4	1	1.08	664.72	1.65
	9	parsley	-0.09	7	2	1.00	688.32	1.70
	10	seaweed	0.26	7	2	1.04	692.24	1.45
	11	spinach	0.21	7	2	1.08	701.24	1.05
	12	turtle	0.61	6	2	1.04	696.12	1.05
Hard Affirmative	1	cabbage	0.57	7	2	1.17	726.29	1.15
	2	cactus	0.12	6	2	1.08	744.63	1.25
	3	dill	-0.05	4	1	1.32	909.56	1.65
	4	fungus	0.20	6	2	1.16	732.96	2.25
	5	grasshopper	0.29	11	3	1.00	710.12	1.10
	6	lime	0.65	4	1	1.04	744.88	1.05
	7	marijuana	0.10	9	4	1.12	886.52	1.45
	8	moss	-0.16	4	1	1.08	750.44	1.55
	9	olive	0.82	5	2	1.24	799.08	1.55
	10	sage	0.13	4	1	1.28	799.68	2.32
	11	spearmint	-0.01	9	2	1.00	768.76	1.65
	12	vine	0.26	4	1	1.16	787.96	1.40
Easy Negative	1	butter	0.39	6	2	2.00	628.60	2.05
	2	cheese	0.32	6	1	2.00	672.96	1.9
	3	cherry	0.80	6	2	2.00	673.67	1.55
	4	cinnamon	0.28	8	3	2.00	738.52	2.3
	5	coffee	0.32	6	2	2.00	614.64	2.15
	6	crab	0.73	4	1	2.00	662.52	1.95
	7	egg	0.24	3	1	1.96	698.20	1.8
	8	nutmeg	-0.01	6	2	2.00	727.64	2.45
	9	rose	0.57	4	1	2.00	722.88	1.65
	10	tiger	0.53	5	2	2.00	672.22	1.7
	11	wheat	0.28	5	1	2.00	673.84	2.1
	12	zebra	0.29	5	2	2.00	632.83	1.8
Hard Negative	1	apricot	0.03	7	3	1.88	743.71	1.55
	2	garlic	0.35	6	2	1.72	841.48	1.85
	3	ginger	0.17	6	2	1.80	839.84	2.35
	4	grape	0.77	5	1	1.72	811.20	1.55
	5	mango	0.50	5	2	1.84	771.80	1.85
	6	mushroom	0.75	8	2	2.00	750.48	1.7
	7	nectarine	-1.02	9	3	1.96	749.04	1.5
	8	panther	0.53	7	2	1.96	771.52	1.9
	9	pepper	0.35	6	2	1.78	767.78	1.75
	10	radish	-0.17	6	2	1.84	804.76	1.5
	11	tick	-0.10	4	1	1.96	794.00	3
	12	tulip	0.01	5	2	1.88	779.50	1.7
Easy Avg			**0.33**	**5.96**	**1.96**	**1.53**	**681.11**	**1.64**
Hard Avg			**0.21**	**6.13**	**1.92**	**1.50**	**782.75**	**1.69**
Modality Avg			**0.27**	**6.04**	**1.94**	**1.51**	**731.93**	**1.66**

Auditory modality (*Sounds Loud?*)

Loud		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Easy Affirmative	1	ambulance	0.24	9	3	1.00	754.74	1.35
	2	bomb	0.86	4	1	1.04	702.59	2.59
	3	dynamite	0.32	8	3	1.04	659.52	2.82
	4	jet	0.72	3	1	1.00	697.20	1.53
	5	missile	-0.13	7	2	1.04	694.33	2.76
	6	motorcycle	0.51	10	4	1.04	711.07	1.25
	7	rocket	0.21	6	2	1.07	715.85	2.56
	8	saxophone	0.09	9	3	1.07	765.85	1.53
	9	shotgun	0.68	7	2	1.00	705.30	2.06
	10	stereo	0.24	6	3	1.15	741.78	1.18
	11	thunder	0.60	7	2	1.00	685.41	1.29
	12	trombone	-0.63	8	2	1.00	708.70	1.76
Hard Affirmative	1	bazooka	-0.66	7	3	1.15	794.78	2.71
	2	blender	0.45	7	2	1.30	865.07	1.53
	3	cannon	0.19	6	2	1.15	776.11	2.56
	4	car	0.85	3	1	1.26	801.22	1.76
	5	drill	0.61	5	1	1.26	798.19	1.65
	6	grenade	-1.18	7	2	1.11	768.26	3.00
	7	guitar	0.46	6	2	1.11	772.52	1.35
	8	hairdryer	0.06	9	3	1.22	803.07	1.41
	9	rifle	-0.16	5	2	1.04	775.69	2.00
	10	train	0.87	5	1	1.11	770.93	1.59
	11	tuba	0.50	4	2	1.00	780.37	1.35
	12	vacuum	0.02	6	2	1.11	788.37	1.24
Easy Negative	1	chair	0.81	5	1	1.96	664.00	2.18
	2	sink	0.70	4	1	2.00	693.00	2.29
	3	pencil	0.90	6	2	1.96	707.04	2.29
	4	ruler	-0.18	5	2	2.00	712.22	2.65
	5	glass	0.32	5	1	1.96	715.50	2.59
	6	rope	0.66	4	1	2.00	717.00	2.35
	7	pen	0.90	3	1	1.96	731.63	1.94
	8	knife	0.80	5	1	1.93	732.93	2.41
	9	stick	0.48	5	1	2.00	745.67	2.35
	10	plate	0.51	5	1	1.96	700.77	2.88
	11	oven	0.70	4	2	1.92	755.00	2.12
	12	scissors	0.36	8	2	1.96	765.15	2.06
Hard Negative	1	pan	0.48	3	1	1.96	771.74	2.24
	2	belt	0.79	4	1	1.96	793.96	2.41
	3	bicycle	0.62	7	3	1.96	794.74	2.00
	4	axe	0.31	3	1	1.81	815.26	2.25
	5	microwave	0.06	9	3	1.70	837.74	1.53
	6	grill	0.56	5	1	1.88	919.58	2.00
	7	refrigerator	-0.10	12	5	1.77	910.92	1.65
	8	clock	0.70	5	1	1.69	865.65	2.06
	9	stream	0.42	6	1	1.92	868.23	2.00
	10	stove	0.62	5	1	1.93	880.48	2.00
	11	scooter	0.24	7	2	1.77	882.15	2.12
	12	wagon	0.34	5	2	1.85	864.56	2.24
Easy Avg			**0.44**	**5.96**	**1.83**	**1.50**	**715.93**	**2.12**
Hard Avg			**0.29**	**5.88**	**1.88**	**1.50**	**820.82**	**1.94**
Modality Avg			**0.37**	**5.92**	**1.85**	**1.50**	**768.37**	**2.03**

Tactile modality (*Feels Soft*?)

Tactile		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Easy Affirmative	1	cheek	0.37	5	1	1.12	787.12	1.29
	2	cheetah	0.04	7	2	1.12	732.73	2.76
	3	cotton	0.28	6	2	1.00	665.04	1.59
	4	fleece	0.32	6	1	1.00	771.89	1.47
	5	lion	0.76	4	2	1.23	730.31	2.41
	6	rabbit	1.04	6	2	1.04	702.74	1.18
	7	satin	0.13	5	2	1.04	688.15	1.18
	8	silk	0.21	4	1	1.00	723.15	1.12
	9	squirrel	0.56	8	2	1.07	769.11	2.24
	10	suede	-0.42	5	1	1.07	751.11	1.41
	11	velvet	0.28	6	2	1.04	685.70	1.41
	12	zebra	0.29	5	2	1.19	771.04	2.59
Hard Affirmative	1	breast	0.67	6	1	1.20	821.24	1.24
	2	cashmere	-0.20	8	2	1.04	810.81	1.41
	3	corduroy	-0.60	8	3	1.40	873.56	1.35
	4	cow	0.70	3	1	1.31	814.31	2.00
	5	felt	0.13	4	1	1.04	816.81	1.35
	6	flannel	0.10	7	2	1.15	819.52	1.29
	7	giraffe	0.21	7	2	1.19	794.48	2.82
	8	goat	0.79	4	1	1.19	809.69	1.88
	9	horse	0.45	5	1	1.19	890.48	1.71
	10	linen	-0.01	5	2	1.19	832.85	1.53
	11	polyester	0.10	9	4	1.31	910.77	1.65
	12	tiger	0.53	5	2	1.19	810.74	2.88
Easy Negative	1	aluminum	0.24	8	4	1.93	684.56	2.82
	2	bronze	0.28	6	1	1.96	696.77	3.06
	3	diamond	0.89	7	3	1.93	710.44	2.65
	4	glass	0.32	5	1	2.00	734.92	1.59
	5	granite	0.24	7	2	2.00	728.74	2.82
	6	marble	0.32	6	2	1.93	705.07	2.47
	7	pebble	-0.04	6	2	1.96	734.04	1.71
	8	platinum	0.17	8	3	1.96	690.85	3.18
	9	rock	0.77	4	1	1.96	713.30	1.65
	10	silver	0.28	6	2	2.00	706.23	2.71
	11	steel	0.35	5	1	2.00	673.22	2.65
	12	tooth	0.79	5	1	1.96	741.37	1.71
Hard Negative	1	bone	0.58	4	1	1.96	748.33	2.53
	2	boulder	-0.58	7	2	2.00	744.44	2.24
	3	cricket	0.61	7	2	2.00	752.63	3.06
	4	crystal	0.39	7	2	1.96	765.74	2.35
	5	elbow	0.39	5	2	1.89	849.44	1.59
	6	elephant	0.55	8	3	1.81	872.35	2.53
	7	forehead	0.35	8	2	1.65	846.88	1.53
	8	hip	0.59	3	1	1.77	811.62	1.47
	9	pearl	0.21	5	1	1.88	827.00	2.12
	10	sapphire	0.06	8	2	2.00	800.26	3.29
	11	shoulder	0.38	9	2	1.81	834.19	1.38
	12	vinyl	0.06	5	2	1.77	926.58	1.24
Easy Avg			**0.32**	**5.87**	**1.79**	**1.52**	**722.06**	**2.06**
Hard Avg			**0.27**	**6.13**	**1.83**	**1.54**	**824.36**	**1.93**
Modality Avg			**0.29**	**6.00**	**1.81**	**1.53**	**773.21**	**2.00**

Gustatory modality (*Tastes Sweet*?)

Sweet		Stimulus	Familiarity	Letters	Syllables	Response	RT	Abstract
Easy Affirmative	1	caramel	0.24	7	3	1.04	718.25	1.10
	2	cherry	0.80	6	2	1.00	642.84	1.50
	3	cookie	0.24	6	2	1.04	677.33	1.25
	4	fudge	0.35	5	1	1.08	658.04	1.30
	5	juice	0.24	5	1	1.04	670.67	1.75
	6	kiwi	-0.73	4	2	1.17	647.08	1.35
	7	lemonade	0.28	8	3	1.16	715.60	1.25
	8	melon	0.57	5	2	1.08	685.17	1.35
	9	papaya	-0.65	6	3	1.39	710.48	2.05
	10	pop	0.13	3	1	1.08	688.60	1.32
	11	raspberry	0.50	9	2	1.17	678.92	1.45
	12	taffy	-0.01	5	2	1.17	660.83	1.35
Hard Affirmative	1	banana	0.67	6	3	1.12	723.04	1.40
	2	cider	0.17	5	2	1.24	778.64	1.20
	3	cocoa	0.13	5	2	1.17	771.00	1.45
	4	coconut	0.52	7	3	1.21	753.92	1.25
	5	cream	0.17	5	1	1.32	739.86	1.55
	6	gum	0.28	3	1	1.08	744.04	1.00
	7	licorice	-0.12	8	3	1.20	837.96	1.55
	8	margarita	0.17	9	4	1.40	894.08	1.65
	9	raisin	0.04	6	2	1.28	758.60	1.25
	10	toffee	-0.38	6	2	1.28	835.36	1.25
	11	truffle	-0.16	7	2	1.25	806.13	1.85
	12	yogurt	0.13	6	2	1.21	728.00	1.45
Easy Negative	1	basil	-0.12	5	2	2.00	616.38	2.45
	2	cabbage	0.57	7	2	1.88	651.21	1.70
	3	carrot	0.93	6	2	1.96	678.72	1.75
	4	garlic	0.35	6	2	1.92	666.33	1.90
	5	lemon	0.77	5	2	1.64	714.32	1.35
	6	lettuce	0.45	7	2	1.92	692.96	1.60
	7	olive	0.82	5	2	1.92	650.76	1.60
	8	parsley	-0.09	7	2	1.96	699.92	2.30
	9	radish	-0.17	6	2	1.96	684.91	2.10
	10	salsa	0.17	5	2	1.92	709.24	1.70
	11	spinach	0.21	7	2	1.96	671.24	1.65
	12	walnut	0.32	6	2	1.88	707.88	1.75
Hard Negative	1	almond	0.24	6	2	1.76	752.96	1.60
	2	butter	0.39	6	2	1.64	794.16	1.75
	3	cashew	0.10	6	2	1.83	808.42	1.50
	4	cheese	0.32	6	1	1.92	763.20	1.60
	5	coffee	0.32	6	2	1.80	803.76	1.95
	6	ginger	0.17	6	2	1.75	797.29	2.00
	7	nutmeg	-0.01	6	2	1.79	848.17	2.05
	8	pecan	-0.09	5	2	1.68	798.72	1.60
	9	pretzel	0.39	7	2	1.79	783.04	1.90
	10	rum	0.43	3	1	1.75	762.17	1.80
	11	squash	0.54	6	1	1.92	734.08	1.85
	12	tomato	0.51	6	3	1.91	746.83	1.75
Easy Avg			**0.26**	**5.88**	**2.00**	**1.51**	**679.07**	**1.62**
Hard Avg			**0.20**	**5.92**	**2.04**	**1.51**	**781.81**	**1.59**
Modality Avg			**0.23**	**5.90**	**2.02**	**1.51**	**730.44**	**1.61**

Experiment 2 items.

<u>Easy Perceptual verification (*Has Four Feet*?)</u>

Four Feet		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Affirmative	1	bear	0.87	4	1	1.00	660.78	1.35
	2	buffalo	0.73	7	3	1.00	669.52	1.47
	3	bull	0.28	4	1	1.00	673.41	1.47
	4	cat	0.94	3	1	1.00	683.70	1.29
	5	cheetah	0.04	7	2	1.00	638.22	1.35
	6	cougar	0.29	6	2	1.04	657.56	1.41
	7	coyote	0.70	6	3	1.00	634.59	1.59
	8	deer	0.46	4	1	1.04	613.30	1.35
	9	dog	1.02	3	1	1.00	628.52	1.29
	10	donkey	0.67	6	2	1.04	703.56	1.41
	11	elephant	0.55	8	3	1.00	701.93	1.35
	12	giraffe	0.21	7	2	1.08	666.54	1.47
	13	goat	0.79	4	1	1.07	645.52	1.35
	14	horse	0.45	5	1	1.04	629.89	1.29
	15	jaguar	0.24	6	2	1.00	649.22	1.41
	16	leopard	0.45	7	2	1.00	660.00	1.47
	17	lion	0.76	4	2	1.00	592.74	1.53
	18	llama	-0.48	5	2	1.04	658.26	1.47
	19	mule	0.66	4	1	1.00	675.56	1.47
	20	panther	0.53	7	2	1.04	639.89	1.53
	21	sheep	1.05	5	1	1.04	640.26	1.35
	22	tiger	0.53	5	2	1.00	649.30	1.35
	23	wolf	0.64	4	1	1.00	649.22	1.35
	24	zebra	0.29	5	2	1.04	650.70	1.35
Negative	1	canary	0.19	6	3	2.00	703.48	2.00
	2	cardinal	0.31	8	3	2.00	718.23	1.71
	3	crow	0.84	4	1	1.96	704.59	1.65
	4	dolphin	0.68	7	2	2.00	650.11	2.65
	5	dove	0.48	4	1	1.93	703.74	1.47
	6	eagle	0.81	5	2	1.96	630.33	1.94
	7	falcon	-0.19	6	2	1.96	688.93	1.88
	8	hawk	0.11	4	1	2.00	721.59	2.00
	9	hummingbird	0.05	11	3	2.00	689.23	2.12
	10	ostrich	-0.26	7	2	2.00	713.85	1.47
	11	owl	0.32	3	1	1.92	671.85	2.24
	12	parakeet	-1.33	8	3	1.96	697.30	1.59
	13	parrot	0.58	6	2	1.96	659.15	2.00
	14	pigeon	0.36	6	2	1.96	694.96	1.94
	15	raven	-0.12	5	2	1.96	706.04	1.71
	16	robin	0.69	5	2	1.93	654.07	1.71
	17	seagull	0.49	7	2	2.00	693.81	2.18
	18	shark	0.82	5	1	1.96	702.81	2.47
	19	sparrow	-0.07	7	2	1.96	718.35	2.24
	20	starfish	0.47	8	2	2.00	694.12	2.94
	21	stork	-0.42	5	1	2.00	701.15	1.82
	22	swordfish	0.18	9	2	2.00	660.63	2.81
	23	trout	0.08	5	1	1.96	679.22	2.24
	24	vulture	-0.21	7	2	2.00	704.81	1.88
Total Avg.			**0.36**	**5.71**	**1.79**	**1.50**	**671.55**	**1.72**

Hard Perceptual Verification (*Has Fur*?)

Fur		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Affirmative	1	antelope	-0.17	8	3	1.15	874.19	1.88
	2	ape	0.44	3	1	1.07	795.15	1.65
	3	beaver	0.52	6	2	1.07	766.15	1.53
	4	boar	-0.05	4	1	1.22	827.44	1.71
	5	buffalo	0.73	7	3	1.11	736.37	1.65
	6	bull	0.28	4	1	1.22	794.85	1.59
	7	chimpanzee	0.22	10	3	1.04	814.00	1.47
	8	donkey	0.67	6	2	1.19	748.96	1.29
	9	elk	-0.01	3	1	1.16	753.68	1.76
	10	gazelle	-0.54	7	2	1.39	834.91	1.59
	11	giraffe	0.21	7	2	1.04	787.19	1.71
	12	goat	0.79	4	1	1.08	763.19	1.29
	13	horse	0.45	5	1	1.19	791.15	1.29
	14	hyena	-1.98	5	3	1.04	763.96	1.65
	15	kangaroo	0.13	8	3	1.07	751.56	1.47
	16	lamb	0.86	4	1	1.00	768.11	1.24
	17	lynx	-0.93	4	1	1.19	807.19	1.75
	18	monkey	0.89	6	2	1.15	738.74	1.47
	19	mouse	0.91	5	1	1.07	779.63	1.35
	20	mule	0.66	4	1	1.11	770.81	1.82
	21	puma	-0.38	4	2	1.12	751.00	1.88
	22	sheep	1.05	5	1	1.04	745.41	1.29
	23	skunk	0.27	5	1	1.11	769.59	1.35
	24	zebra	0.29	5	2	1.07	744.63	1.53
Negative	1	canary	0.19	6	3	1.96	762.30	2.13
	2	cardinal	0.31	8	3	1.96	708.41	1.94
	3	chicken	0.87	7	2	1.85	753.26	1.82
	4	crocodile	0.27	9	3	2.00	717.11	2.59
	5	dove	0.48	4	1	1.96	738.85	2.00
	6	duck	0.76	4	1	1.93	744.41	1.65
	7	eagle	0.81	5	2	1.96	773.26	1.59
	8	elephant	0.55	8	3	1.77	858.69	1.76
	9	falcon	-0.19	6	2	1.96	821.70	2.12
	10	goose	0.20	5	1	1.93	739.85	1.82
	11	grasshopper	0.29	11	3	2.00	724.74	2.41
	12	hen	0.27	3	1	1.92	775.69	1.59
	13	hippopotamus	-0.30	12	5	1.81	876.33	2.00
	14	lizard	0.69	6	2	2.00	718.33	2.06
	15	owl	0.32	3	1	1.88	758.92	1.76
	16	parrot	0.58	6	2	1.96	702.19	1.76
	17	penguin	0.55	7	2	1.96	716.33	1.94
	18	pig	0.92	3	1	1.67	829.52	1.82
	19	robin	0.69	5	2	1.89	773.89	2.06
	20	rooster	0.39	7	2	1.93	715.15	1.71
	21	seagull	0.49	7	2	2.00	748.67	2.00
	22	swan	-0.22	4	1	1.96	754.63	1.76
	23	turkey	0.88	6	2	2.00	829.11	1.88
	24	vulture	-0.21	7	2	1.93	715.00	2.06
Total Avg.			**0.31**	**5.79**	**1.88**	**1.52**	**769.46**	**1.74**

Easy Abstract Verification (*Can be Trained*?)

Train		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Affirmative	1	ape	0.44	3	1	1.16	737.84	2.45
	2	cat	0.94	3	1	1.17	715.71	2.25
	3	cheetah	0.04	7	2	1.40	733.04	3.60
	4	chimpanzee	0.22	10	3	1.17	682.63	2.47
	5	cougar	0.29	6	2	1.32	755.32	3.60
	6	cow	0.70	3	1	1.39	782.17	3.20
	7	dog	1.02	3	1	1.08	578.63	2.25
	8	dolphin	0.68	7	2	1.12	739.96	2.75
	9	donkey	0.67	6	2	1.32	762.88	3.05
	10	elephant	0.55	8	3	1.32	805.96	3.30
	11	goat	0.79	4	1	1.33	718.00	3.20
	12	gorilla	0.24	7	3	1.28	714.24	2.55
	13	horse	0.45	5	1	1.04	696.35	2.60
	14	kangaroo	0.13	8	3	1.40	735.76	3.45
	15	leopard	0.45	7	2	1.32	743.96	3.30
	16	lion	0.76	4	2	1.32	683.56	2.70
	17	monkey	0.89	6	2	1.09	702.39	2.45
	18	mule	0.66	4	1	1.24	716.28	3.45
	19	parakeet	-1.33	8	3	1.32	809.80	2.85
	20	parrot	0.58	6	2	1.22	684.74	2.65
	21	pig	0.92	3	1	1.36	784.32	3.05
	22	rabbit	1.04	6	2	1.25	701.88	3.30
	23	tiger	0.53	5	2	1.28	720.96	2.65
	24	whale	0.77	5	1	1.36	804.12	3.55
Negative	1	ant	0.73	3	1	1.88	621.72	4.30
	2	antelope	-0.17	8	3	1.63	702.54	3.45
	3	beetle	0.54	6	2	1.88	682.44	4.10
	4	centipede	-0.53	9	3	1.87	703.52	4.32
	5	cockroach	-0.25	9	2	1.88	677.00	4.40
	6	coyote	0.70	6	3	1.65	684.78	3.70
	7	crab	0.73	4	1	1.91	703.91	4.10
	8	flea	0.39	4	1	1.84	672.88	4.10
	9	hornet	0.31	6	2	1.88	660.20	4.30
	10	maggot	0.20	6	2	1.96	674.96	4.20
	11	mosquito	0.27	8	3	1.92	688.75	4.25
	12	moth	-0.23	4	1	1.88	696.75	4.25
	13	pigeon	0.36	6	2	1.61	668.04	3.40
	14	raccoon	0.27	7	2	1.61	696.00	3.70
	15	salmon	0.41	6	2	1.83	710.29	4.25
	16	shark	0.82	5	1	1.75	713.42	3.45
	17	shrimp	0.94	6	1	1.88	691.60	4.45
	18	spider	0.56	6	2	1.75	703.83	4.10
	19	starfish	0.47	8	2	1.84	664.08	4.30
	20	termite	0.20	7	2	1.92	696.24	4.20
	21	tick	-0.10	4	1	1.83	680.75	4.40
	22	trout	0.08	5	1	1.88	658.50	4.15
	23	tuna	0.85	4	2	1.83	680.75	4.30
	24	wasp	-0.01	4	1	1.96	698.30	4.35
Total Avg.			**0.38**	**5.80**	**1.86**	**1.55**	**710.63**	**3.54**

Hard Abstract Verification (*Lays Eggs*?)

Eggs		**Stimulus**	**Familiarity**	**Letters**	**Syllables**	**Response**	**RT**	**Abstract**
Affirmative	1	alligator	0.50	9	4	1.22	838.04	3.06
	2	ant	0.73	3	1	1.33	853.96	3.82
	3	bee	0.73	3	1	1.23	843.08	3.65
	4	cricket	0.61	7	2	1.22	807.78	4.00
	5	flea	0.39	4	1	1.19	768.41	3.41
	6	hornet	0.31	6	2	1.27	820.31	3.65
	7	jellyfish	0.60	9	2	1.35	880.35	3.76
	8	ladybug	0.18	7	3	1.24	783.92	3.65
	9	lizard	0.69	6	2	1.15	844.85	3.29
	10	lobster	0.68	7	2	1.19	961.93	3.76
	11	mackerel	-1.23	8	3	1.24	910.88	4.00
	12	mosquito	0.27	8	3	1.19	881.44	3.71
	13	moth	-0.23	4	1	1.22	874.96	3.65
	14	octopus	0.26	7	3	1.38	991.42	3.53
	15	owl	0.32	3	1	1.38	780.42	3.18
	16	salamander	-0.69	10	4	1.19	893.33	3.47
	17	shark	0.82	5	1	1.36	829.12	3.53
	18	shrimp	0.94	6	1	1.19	870.27	3.82
	19	spider	0.56	6	2	1.15	853.19	3.29
	20	starfish	0.47	8	2	1.21	834.67	3.47
	21	stingray	-0.97	8	2	1.27	1009.96	3.93
	22	termite	0.20	7	2	1.19	819.65	4.12
	23	tuna	0.85	4	2	1.15	792.08	3.41
	24	wasp	-0.01	4	1	1.22	867.93	3.35
Negative	1	bear	0.87	4	1	1.96	771.19	3.41
	2	beaver	0.52	6	2	1.93	767.04	3.12
	3	boar	-0.05	4	1	1.96	755.50	3.18
	4	buffalo	0.73	7	3	1.93	752.07	3.24
	5	cheetah	0.04	7	2	1.88	745.72	3.12
	6	chipmunk	0.11	8	2	1.88	856.08	3.47
	7	dog	1.02	3	1	1.93	738.48	3.35
	8	dolphin	0.68	7	2	1.74	880.89	3.82
	9	elk	-0.01	3	1	1.92	767.50	3.12
	10	gazelle	-0.54	7	2	1.81	807.52	3.29
	11	jaguar	0.24	6	2	1.96	768.78	3.59
	12	lamb	0.86	4	1	1.89	767.78	3.12
	13	leopard	0.45	7	2	1.93	741.07	3.29
	14	llama	-0.48	5	2	1.92	762.04	3.24
	15	lynx	-0.93	4	1	1.93	775.19	3.12
	16	mouse	0.91	5	1	1.92	870.96	3.24
	17	mule	0.66	4	1	1.96	768.56	3.24
	18	puma	-0.38	4	2	1.93	812.33	3.35
	19	raccoon	0.27	7	2	1.96	784.52	3.53
	20	skunk	0.27	5	1	1.89	725.19	3.65
	21	squirrel	0.56	8	2	1.81	800.73	3.35
	22	tiger	0.53	5	2	1.96	731.33	3.35
	23	whale	0.77	5	1	1.78	917.56	3.71
	24	zebra	0.29	5	2	1.96	814.70	3.29
Total Avg.			**0.30**	**5.81**	**1.81**	**1.57**	**822.81**	**3.47**

BIBLIOGRAPHY

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

Armstrong, S. L., Gleitman, L.R., Gleitman, H. (1983). What some concepts might not be. *Cognition, 13*, 263-308.

Barde, L. H. F., & Thompson-Schill, S. L. (2002). Models of functional organization of the lateral prefrontal cortex in verbal working memory: Evidence in favor of the process model. *Journal of Cognitive Neuroscience*, 14, 1054-1063.

Barsalou, L.W. (1991). Deriving categories to achieve goals. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, pp. 1-64). San Diego, CA: Academic Press.

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

Barsalou, L.W. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, 358, 1177-1187.

Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2002). Parallel visual motion processing streams for manipulable objects and human movements. *Neuron, 34*, 149-159.

Bertenthal, B. I. (1993). Infant's perception of biomechanical motions: Intrinsic image and knowledge-based constraints. In C. Granrud (Ed.), *Visual perception and cognition in infancy.* Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, *94*, 115–147.

Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S., Rao, S. M., & Cox, R. W. (1999). Conceptual processing during the conscious resting state. A functional MRI study. *Journal of Cognitive Neuroscience, 11*, 80-95.

Binder, J. R., Frost, J. A., Hammeke, T. A., Cox, R. W., Rao, S. M., & Prieto, T. (1997). Human brain language areas identified by functional magnetic resonance imaging. *Journal of Neuroscience, 17*, 353-362.

Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition, 60*, 1-29.

Bloom, P. (1998). Theories of artifact categorization. *Cognition*, 66, 87-93.

Bookheimer, S. (2002). Functional MRI of Language: New Approaches to Understanding the Cortical Organization of Semantic Processing. *Annual Review of Neuroscience, 25*, 151-188.

Borgo, F., & Shallice, T. (2001). When living things and other 'sensory quality' categories behave in the same fashion: a novel category specificity effect. *Neurocase*, *7*, 201-220.

Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human v1. *Journal of Neuroscience*, *16*, 4207-4221.

Brewer, J. B., Zhao, Z., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. (1998). Making memories: brain activity that predicts how well visual experience will be remembered. *Science, 281*, 1185-1187.

Buchel, C., Price, C., & Friston, K. (1998). A multimodal language region in the ventral visual pathway. *Nature, 394*, 274-277.

Buckner, R. L., Bandettini, P. A., O'Craven, K. M., Savoy, R. L., Petersen, S. E., Raichle, M. E., & Rosen, B. R. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences of the United States of America, 93,* 14878-14883.

Buxbaum, L. J., & Saffran, E. M. (2002). Knowledge of object manipulation and object function: dissociations in apraxic and nonapraxic subjects. *Brain and Language, 82*, 179-199.

Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience, 12*, 1-47.

Cardillo, E.R., Aydelott, J., Matthews, P.M., & Devlin, J.T. (2004). Left inferior prefrontal cortex activity reflects inhibitory rather than facilitatory priming. *Journal of Cognitive Neuroscience, 16*, 1552-1561.

Caramazza, A. (1996). Neuropsychology: The brain's dictionary. *Nature, 380*, 485-486.

Caramazza, A. (1998). The interpretation of semantic category-specific deficits: What do they reveal about the organization of conceptual knowledge in the brain? *Neurocase, 4*, 265-272.

Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, 7(3), 161-189.

Caramazza, A. & Mahon, B. Z. (2003). The organization of conceptual knowledge: the evidence from category-specific semantic deficits. *Trends in Cognitive Sciences 7*, 354-361.

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

Carey, S. (1985). *Conceptual change in childhood*. Cambridge: MIT Press.

Chao, L. L., & Martin, A. (1999). Cortical regions associated with perceiving, naming, and knowing about colors. *Journal of Cognitive Neuroscience, 11*, 25-35.

Chao, L. L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage, 12*, 478-484.

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

Chao, L. L., Weisberg, J., & Martin, A. (2002). Experience-dependent modulation of category-related cortical activity. *Cerebral Cortex, 12*, 545-551.

Chee, M., Hon, N., Caplan, D., Lee, H. L., & Goh, J. (2002). Frequency of concrete words modulates prefrontal activation during semantic judgments. *Neuroimage, 16*, 259-268.

Collins, A. M., & Loftus, E.L. (1975). A spreading activation theory of semantic processing. *Psychological Review, 84*, 407-428.

Coltheart, M. (2001). Assumptions and methods in cognitive neuropsychoplogy. In B. Rapp (Ed.), *The handbook of cognitive neuropsychology*: *What deficits reveal about the human mind*. Philadelphia, PA: Psychology Press.

Dale, A. M., Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping, 5,* 329-340.

Damasio, H., Grabowski, T. J., Tranel, D., Hichwa, R. D., & Damasio, A. R. (1996). A neural basis for lexical retrieval. *Nature, 380*, 499-505.

De Renzi, E., & Lucchelli, F. (1994). Are semantic systems separately represented in the brain? The case of living category impairment. *Cortex, 30*, 3-25.

Demb, J. B., Desmond, J. E., Wagner, A. D., Vaidya, C. J., Glover, G. H., & Gabrieli, J. D. (1995). Semantic encoding and retrieval in the left inferior prefrontal cortex: a functional MRI study of task difficulty and process specificity. *Journal of Neuroscience, 15*, 5870-5878.

Demonet, J. F., Chollet, F., Ramsay, S., Cardebat, D., Nespoulous, J. L., Wise, R., Rascol, A., & Frackowiak, R. (1992). The anatomy of phonological and semantic processing in normal subjects. *Brain, 115*, 1753-1768.

Dennett, D. C. (1984). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.

Devlin, J. T., Gonnerman, L. M., Andersen, E. S., & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: a computational account. *Journal of Cognitive Neuroscience, 10*, 77-94.

Devlin, J. T., Matthews, P. M., & Rushworth, M. F. S. (2003). Semantic processing in the left inferior prefrontal cortex: A combined functional magnetic resonance imaging and transcranial magnetic stimulation study. *Journal of Cognitive Neuroscience, 15*, 71-84.

Devlin, J. T., Moore, C. J., Mummery, C. J., Gorno Tempini, M. L., Phillips, J. A., Noppeney, U., Frackowiak, R. S. J., Friston, K. J., & Price, C. J. (2002a). Anatomic constraints on cognitive theories of category specificity. *Neuroimage, 15*, 675-685.

Devlin, J. T., Russell, R. P., Davis, M. H., Price, C. J., Moss, H. E., Fadili, M. J., & Tyler, L. K. (2002b). Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. *Neuropsychologia, 40*, 54-75.

Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: modality specificity and emergent category specificity. *Journal of Experimental Psychology, General, 120*, 339-357.

Fiez, J. A. (1996). Cerebellar contributions to cognition. *Neuron, 16*, 13-15.

Fiez, J. A. (1997). Phonology, semantics, and the role of the left inferior prefrontal cortex. *Human Brain Mapping, 5*, 79-83.

Forde, E. M. E., Francis, D., Riddoch, M. J., Rumiati, R. I., & Humphreys, G. W. (1997). On the links between visual knowledge and naming: A single case study of a patient with a category-specific impairment for living things. *Cognitive Neuropsychology, 14*, 403-458.

Fuster, J. M. (1995). *Memory in the cerebral cortex*. Cambridge, MA: MIT Press.

Gabrieli, J. D., Poldrack, R. A., & Desmond, J. E. (1998). The role of left prefrontal cortex in language and memory. Proceedings of the National Academy of Sciences of the United States of America, 95, 906-913.

Gainotti, G., & Silveri, M. C. (1996). Cognitive and anatomical locus of lesion in a patient with a category-specific semantic impairment for living beings. *Cognitive Neuropsychology, 13*, 357-389.

Gelman, S. A., & Opfer, J. E. (2002). Development of the animate-inanimate distinction. In U. Goswami (Ed.), *Blackwell Handbook of Childhood Cognitive Development* (pp. 151 - 166). Oxford, UK: Blackwell.

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

Gold, B. T. and R. L. Buckner (2002). Common prefrontal regions coactivate with dissociable posterior regions during controlled semantic and phonological tasks. *Neuron, 35,* 803-812.

Goldberg, R. F. (2001). Does adult category verification reflect child-like concepts? Paper presented at the Cognitive Science Society meeting, Edinburgh, Scotland.

Goldberg, R. F., Perfetti, C. A., Schneider, W. (2003). What's 'alive' in semantic knowledge? Cognitive Neuroscience Society, New York, NY.

Goldberg, R. F., Suzuki, I., Perfetti, C. A., & Schneider, W. (2002). Deriving category structure within semantic memory from similarity-based relations using neuroimaging. Society for Neuroscience, Orlando, FL.

Goldstone, R. L. (1994). The role of similarity in categorization: providing a groundwork. *Cognition, 52*, 125-157.

Goldstone, R. L., & Barsalou, L. W. (1998). Reuniting perception and conception. *Cognition, 65*, 231-262.

Hahn, U. C., & Chater, N. (1998). Similarity and rules: Distinct? Exhaustive? Empirically distinguishable? *Cognition, 65*, 197-230.

Hampton, J. A. (1997). Associative and similarity-based process in categorization decisions. *Memory & Cognition, 25*, 625-640.

Hart, J., Jr., & Gordon, B. (1992). Neural subsystems for object knowledge. *Nature, 359*, 60-64.

Hart, J., Jr., Berndt, R. S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature, 316*, 439-440.

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

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science, 293*, 2425-2430.

Hills, A. E., & Caramazza, A. (1991). Mechanisms for accessing lexical representations for output: Evidence from a category-specific semantic deficit. *Brain and Language, 40*, 106-144.

Hodges, J. R., Bozeat, S., Lambon Ralph, M. A., Patterson, K., & Spatt, J. (2000). The role of conceptual knowledge in object use evidence from semantic dementia. *Brain, 123*, 1913-1925.

Humphreys, G. W., & Forde, E. M. (2001). Hierarchies, similarity, and interactivity in object recognition: "category-specific" neuropsychological deficits. *Behavioral and Brain Sciences, 24*, 453-476.

James, T. W. and I. Gauthier (2003). Auditory and action semantic features activate sensory-specific perceptual brain regions. *Current Biology, 13*, 1792-1796.

Kapur, S., Rose, R., Liddle, P. F., Zipursky, R. B., Brown, G. M., Stuss, D., Houle, S., & Tulving, E. (1994). The role of the left prefrontal cortex in verbal processing: semantic processing or willed action? *Neuroreport, 5*, 2193-2196.

Katz, J., & Fodor, J.A. (1963). The structure of a semantic theory. *Language, 39*, 170-210.

Kerns, J.G., Cohen, J.D., Stenger, V.A., Carter, C.S. (2004). Prefrontal cortex guides context-appropriate responding during language production. *Neuron, 43*, 283–291.

Keil, F. C. (1989). *Concepts, Kinds, and Cognitive Development*. Cambridge: MIT Press.

Kellenbach, M. L., Brett, M., & Patterson, K. (2001). Large, colorful, or noisy? Attribute- and modality-specific activations during retrieval of perceptual attribute knowledge. *Cognitive, Affective, & Behavioral Neuroscience, 1*, 207-221.

Lambon Ralph, M. A., Graham, K. S., Patterson, K., & Hodges, J. R. (1999). Is a picture worth a thousand words? Evidence from concept definitions by patients with semantic dementia. *Brain and Language, 70*, 309-335.

Le Clec'H, G., Dehaene, S., Cohen, L., Mehler, J., Dupoux, E., Poline, J. B., Lehericy, S., van de Moortele, P. F., & Le Bihan, D. (2000). Distinct Cortical Areas for Names of Numbers and Body Parts Independent of Language and Input Modality. *Neuroimage, 12*, 381-391.

Lloyd-Jones, T. J., & Luckhurst, L. (2002). Outline shape is a mediator of object recognition that is particularly important for living things. *Memory & Cognition, 30*, 489-498.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature, 412*, 150-157.

Malt, B. C., & Johnson, E. C. (1998). Artifact category membership and the intentional-historical theory. *Cognition, 66*, 79-85.

Marques, J. F. (2002). An attribute is worth more than a category: Testing different semantic memory organization hypotheses in relation to the living/nonlivng things dissociation. *Cognitive Neuropsycholgy, 19*, 463-478.

Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science, 270*, 102-105.

Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature, 379*, 649-652.

Matsumoto, K., Suzuki, W., & Tanaka, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science, 301*, 229-232.

McCarthy, R. A., & Warrington, E. K. (1988). Evidence for modality-specific meaning systems in the brain. *Nature, 334*, 428-430.

McCloskey, M. & Glucksberg, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology, 11*, 1-37.

McDermott, K. B., Buckner, R. L., Petersen, S. E., Kelley, W. M., & Sanders, A. L. (1999). Set- and code-specific activation in frontal cortex: an fMRI study of encoding and retrieval of faces and words. *Journal of Cognitive Neuroscience, 11*, 631-640.

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

Middleton, F., & Strick, P. L. (1994). Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science, 266*, 458-461.

Miller, E. K., Freedman, D. J., & Wallis, J. D. (2002). The prefrontal cortex: categories, concepts, and cognition. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences, 357*, 1123-1136.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167-202.

Moore, C. J., & Price, C. J. (1999). A functional neuroimaging study of the variables that generate category-specific object processing differences. *Brain, 122*, 943-962.

Moss, H. E., & Tyler, L. K. (2000). A progressive category-specific semantic deficit for non-living things. *Neuropsychologia, 38*, 60-82.

Moss, H. E., Tyler, L. K., & Jennings, F. (1997). When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology, 14*, 901-950.

Moss, H. E., Tyler, L. K., Durrant Peatfield, M., & Bunn, E. M. (1998). Two eyes of a see-through: Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase, 4*, 291-310.

Mummery, C. J., Patterson, K., Hodges, J. R., & Price, C. J. (1998). Functional neuroanatomy of the semantic system: divisible by what? *Journal of Cognitive Neuroscience, 10*, 766-777.

Mummery, C. J., Patterson, K., Hodges, J. R., & Wise, R. J. (1996). Generating 'tiger' as an animal name or a word beginning with T: differences in brain activation. *Proceedings of the Royal Society of London Series B: Biological sciences, 263*, 989-995.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92,* 289-316.

Noppeney, U., Phillips, J., & Price, C. (2004). The neural areas that control the processing and selection of semantics. *Neuropsychologia, 42*, 1269-1280.

Noppeney, U., & Price, C. (2002). Retrieval of Visual, Auditory, and Abstract Semantics. *Neuroimage, 15*, 917-926.

Oliver, R. T. and S. L. Thompson-Schill (2003). Dorsal stream activation during retrieval of object size and shape. *Cognitive, Affective, & Behavioral Neuroscience, 3*, 309-322.

Patalano, A. L., Smith, E. E., Jonides, J., & Koeppe, R. (2001). PET evidence for multiple strategies for categorization. *Cognitive, Affective, & Behavioral Neuroscience, 1*, 360-370.

Perani, D., Cappa, S. F., Bettinardi, V., Bressi, S., Gorno Tempini, M., Matarrese, M., & Fazio, F. (1995). Different neural systems for the recognition of animals and man-made tools. *Neuroreport, 6*, 1637-1641.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature, 331*, 585-589.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1989). Positron emission tomographic studies of the processing of single words. *Journal of Cognitive Neuroscience, 1*, 153-170.

Pietrini, V., Nertempi, P., Vaglia, A., Revello, M. G., Pinna, V., & Ferro Milone, F. (1988). Recovery from herpes simplex encephalitis: selective impairment of specific semantic categories with neuroradiological correlation. *Journal of Neurology, Neurosurgery, and Psychiatry, 51*, 1284-1293.

Posner, M. I., Petersen, S. E., Fox, P. T., & Raichle, M. E. (1988). Localization of cognitive operations in the human brain. *Science, 240*, 1627-1631.

Price, C. J., Moore, C. J, Humphreys, G. W., & Wise, R. J. S. (1997). Segregating semantic from phonological processes during reading. *Journal of Cognitive Neuroscience, 9*, 727-733.

Price, C.J., & Humphreys, G. W. (1989). The effects of surface detail on object categorisation and naming. *Quarterly Journal of Experimental Psychology, 41A*, 797-828.

Price, C. J., Wise, R. J., & Frackowiak, R. S. (1996). Demonstrating the implicit processing of visually presented words and pseudowords. *Cerebral Cortex, 6*, 62-70.

Rips, L. J. (1995). The current status of research on concept combination. *Mind & Language, 10*, 72-104.

Rips, L. J., Shoben, E.J., Smith, E.E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior, 1*, 1-20.

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

Rolls, E. T. (2000). Memory systems in the brain. *Annual Review of Psychology, 51*, 599-630.

Rolls, E. T. (2001). The rules of formation of the olfactory representations found in the orbitofrontal cortex olfactory areas in primates. *Chemical Senses, 26*, 595-604.

Romanski L. M., Goldman-Rakic, P. S. (2002). An auditory domain in primate prefrontal cortex. *Nature Neuroscience, 5*, 15-16.

Rosch, E. M., Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7,* 573-605.

Roskies, A. L., Fiez, J. A., Balota, D. A., Raichle, M. E., & Petersen, S. E. (2001). Task-dependent modulation of regions in the left inferior frontal cortex during semantic processing. *Journal of Cognitive Neuroscience, 13*, 829-843.

Sacchett, C., & Humphreys, G. W. (1992). Calling a squirrel a squirrel but a canoe a wigwam: A category-specific deficit for artefactual objects and body parts. *Cognitive Neuropsychology, 9*, 73-86.

Saffran, E. M., & Schwartz, M. F. (1994). Of Cabbages and Things: Semantic Memory from a Neuropsychological Perspective--A Tutorial Review. In C. Umilta & M. Moscovitch (Eds.), *Attention & Performance XV* (pp. 507-536). Cambridge, MA: MIT Press.

Schneider, W. & Chein, J. M. (2003). Controlled and automatic processing: behavior, theory, and biological mechanisms. *Cognitive Science, 27*, 525-559.

Sergent, J., Ohta, S., MacDonald, B., & Pulvermueller, F. (1992). Functional neuroanatomy of face and object processing. A positron emission tomography study. *Brain, 1*, 15-36.

Shelton, J. R., & Caramazza, A. (2001). The organization of semantic memory. In B. Rapp (Ed.), *The handbook of cognitive neuropsychology: What deficits reveal about the human mind* (pp. 423-443). Philadelphia, PA: Psychology Press.

Shelton, J. R., Fouch, E., & Caramazza, A. (1998). The selective sparing of body part knowledge: A case study. Neurocase, 4, 339-351.

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

Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65*, 167-196.

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

Talairach, J. and P. Tournoux (1988). *Co-planar stereotaxic atlas of the human brain*. New York, Thieme Medical Publishers, Inc.

Tan, L. H., Liu, H.-L., Perfetti, C. A., Spinks, J. A., Fox, P. T., & Gao, J.-H. (2001). The neural system underlying Chinese logograph reading. *NeuroImage, 13*, 836–846.

Thompson-Schill, S. L., Aguirre, G. K., D'Esposito, M., & Farah, M. J. (1999). A neural basis for category and modality specificity of semantic knowledge. *Neuropsychologia, 37*, 671-676.

Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proceedings of the National Academy of Sciences of the United States of America, 94*, 14792-14797.

Thompson-Schill, S. L., D'Esposito, M., & Kan, I. P. (1999). Effects of repetition and competition on activity in left prefrontal cortex during word generation. *Neuron, 23*, 513-522.

Toglia, M.P. & Battig, W.F. (1978). *Handbook of Semantic Word Norms.* Hillsdale, NJ: Erlbaum.

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

Tyler, L. K., Russell, R., Fadili, J., & Moss, H. E. (2001). The neural representation of nouns and verbs: PET studies. *Brain, 124*, 1619-1634.

Ungerleider, L. G., & Mishkin, M. (1992). Two cortical visual systems. In D. J. Ingle & M. A. Goodale & R. J. W. Mansfield (Eds.*), Analysis of visual behavior* (pp. 549-586). Cambridge, MA: MIT Press.

Vandenberghe, R., Nobre, A. C., & Price, C. (2002). The Response of Left Temporal Cortex to Sentences. *Journal Cognitive Neuroscience, 14*, 550-560.

Vandenberghe, R., Price, C. J., Wise, R., Josephs, O., & Frackowiak, R. S. (1996). Functional anatomy of a common semantic system for words and pictures. *Nature, 383*, 254-256.

Wagner, A. D., Desmond, J. E., Demb, J. B., Glover, G. H., & Gabrieli, J. D. E. (1997). Semantic repetition priming for verbal and pictorial knowledge: A functional MRI study of left inferior prefrontal cortex. *Journal of Cognitive Neuroscience, 9*, 714-726.

Wagner, A. D., Koustaal, W., Maril, A., Schacter, D., & Buckner, R. L. (2000). Task-specific repetition priming in left inferior prefrontal cortex. *Cerebral Cortex, 10*, 1176-1184.

Wagner, A. D., Pare-Blagoev, E. J., Clark, J., & Poldrack, R. A. (2001). Recovering meaning: Left prefrontal cortex guides controlled semantic retrieval. *Neuron, 31*, 329-338.

Wagner, A. D., Schacter, D. L., Rotte, M., Koutstaal, W., Maril, A., Dale, A. M., Rosen, B. R., & Buckner, R. L. (1998). Building memories: remembering and forgetting of verbal experiences as predicted by brain activity. *Science, 281*, 1188-1191.

Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411, 953-956.

Warrington, E. K., & McCarthy, R. (1983). Category specific access dysphasia. *Brain, 106*, 859-878.

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

Warrington, E. K., & McCarthy, R. A. (1994). Multiple meaning systems in the brain: a case for visual semantics. *Neuropsychologia, 32*, 1465-1473.

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

Wiemer-Hastings, K., Krug, J., & Xu, X. (2001) Imagery, context availability, contextual constraint, and abstractness. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pp. 1134–1139. Mahwah, NJ: Lawrence Erlbaum.

Wilson, F. A., Scalaidhe, S. P., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science, 260*, 1955-1958.

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