



Published in final edited form as:

Cogn Neuropsychol. 2012 ; 29(3): 237–248. doi:10.1080/02643294.2012.706218.

Conceptual Representations of Perceptual Knowledge

Edward E. Smith^{1,2}, Nicholas Myers³, Umrao Sethi⁴, Spiro Pantazatos⁵, Ted Yanagihara⁵, and Joy Hirsch⁵

¹Department of Psychology, Columbia University

²Department of Cognitive Neuroscience, New York State Psychiatric Institute

³Department of Experimental Psychology, University of Oxford

⁴Department of Philosophy, University of California, Berkeley

⁵Program for Imaging and Cognitive Science, Columbia University

Abstract

Many neuroimaging studies of semantic memory have argued that knowledge of an object's perceptual properties are represented in a modality-specific manner. These studies often base their argument on finding activation in the left-hemisphere fusiform gyrus - a region assumed to be involved in perceptual processing - when the participant is verifying verbal statements about objects and properties. In this paper we report an extension of one of these influential papers—Kan, Barsalou, Solomon, Minor, and Thompson-Schill (2003)—and present evidence for an amodal component in the representation and processing of perceptual knowledge. Participants were required to verify object-property statements (e.g., “cat- whiskers?”; “bear-wings?”) while they were being scanned by fMRI. We replicated Kan et al's activation in the left fusiform gyrus, but also found activation in regions of left inferior frontal gyrus (IFG) and middle-temporal gyrus, areas known to reflect amodal processes or representations. Further, only activations in the left IFG, an amodal area, were correlated with measures of behavioral performance.

Keywords

Perceptual knowledge; representations; semantics; neuroimaging; modality-specificity

Introduction

Behavioral studies of semantic memory in the 1970s and 1980s assumed that our knowledge of the perceptual properties of everyday objects are represented by abstract or amodal symbols (e.g., Collins & Loftus, 1975; Miller & Johnson-Laird, 1976; Smith, Shoben, & Rips, 1974; Smith & Medin, 1981). That assumption was seriously challenged by the advent of neuroimaging studies of semantic memory in the 1990s. The latter studies produced results that suggested that knowledge about perceptual properties are represented in the relevant modality, visual features in the visual modality, auditory features in the auditory modality, and so on. In a seminal study, Martin, Haxby, Lalonde, Wiggs, and Ungerleider (1995) asked participants to verify simple verbal statements about the colors of common objects -- “A banana is yellow?” -- and found that verification was accompanied by activation in a region of the brain known to be involved in the perception of color, namely the left hemisphere inferior fusiform gyrus in ventral-temporal cortex (e.g., Beauchamp,

Haxby, Jennings, & DeYoe, 1999). Subsequent studies that also required the verification of perceptual properties found similar results (e.g., Chao, Haxby, & Martin, 1999; Mummery, Patterson, Hodges, & Price, 1998; Kellenbach, Brett, & Patterson, 2001; Lee et al., 2002). And one of the best-known reviews of neuroimaging studies of semantic-memory concluded that "...semantic memory is not amodal: each attribute-specific system is tied to a sensorimotor modality (e.g., vision) and even to a specific property within that modality (e.g., color). Information about each feature of a concept is stored within the processing streams that were active during the acquisition of that feature" (Thompson-Schill, 2003, p. 283). Since this review other property-verification experiments have provided further evidence that our knowledge of objects is "perceptually embodied" (Kan et al., 2003; Simmons et al., 2007), as verbal questions consistently activate visual (or motor) areas (e.g., Barsalou, Simmons, Barbey, & Wilson, 2003; Goldberg, Perfetti, & Schneider, 2006; Hauk, Johnsrude & Pulvermuller, 2004; Simmons, Ramjee, McRae, Martin, & Barsalou, 2006; Simmons et al., 2007). (Interestingly, Allport, 1985, articulated this perceptual embodiment hypothesis prior to the advent of the use of neuroimaging in this area.)

In this paper we employ a property-verification paradigm *identical* to one employed in some of the above studies, and report neuroimaging results that raise concerns about the perceptual-embodiment hypothesis. We show that: (1) neural areas associated with *amodal* processes and representations are also activated when people verify statements about perceptual properties; and (2) some of these activations correlate with behavioral performance in the verification task. Finding (1) is not novel – some of the above cited papers found activations in frontal regions that are thought to mediate amodal functions, though these papers emphasized the activations in the more perceptual regions. Findings (1) and (2) together, though, are more informative as they implicate a role for amodal processes or amodal representations of perceptual information.

A comment is in order about what we mean by an "amodal representation". Beyond the literal meaning--that the representation cannot be tied to a particular sensory modality--we mean that the representation is *abstract* in the sense that it cannot be put in a one-to-one correspondence with a perceptual event (for discussion, see, e.g., Kosslyn, 1980). What conditions foster abstract representations? Even when knowledge is initially acquired by direct, perceptual experience, it may eventually come to be represented abstractly; examples would be the knowledge that "cars have brakes" or that "tigers have skin". But an even stronger case can be made for the plausibility of the abstract representations when the knowledge being represented has likely been acquired by verbal communication--e.g., "lions have ribs" (Goldberg et al., 2007). The above examples have all been used in property-verification studies.

Our experiment is based on a prior study by Kan, Barsalou, Solomon, Minor, & Thompson-Schill (2003). This experiment has a couple of advantages. First, the items used by Kan et al. (2003) were selected with unusual care such that the true and false properties associated with a particular concept were equally associated with that concept; this insured that verification had to be based on meaning, not simple associations. (Indeed, the main point of Kan et al.'s study was to show that only when verification is based on meaning is there activation in the fusiform gyrus. We do not dispute this important conclusion.) Second, the items included many kinds of perceptual properties, rather than just 2 or 3 perceptual attributes (e.g., color, size) that participants are continuously asked about (e.g., "Respond positively to any of the following objects that are red in color"). The latter kind of study may lead to a task-specific strategy, such as the participant trying to think of red objects in advance, whereas the use of different properties in different sentences as in Kan et al. (2003) seems closer to natural situations in which people retrieve knowledge about objects.

In the experimental condition of interest in Kan et al. (2003), participants were presented items that contained a category word on top of a screen (e.g., “cat”), and a possible property of that category on the bottom (e.g., “whiskers”). Participants had to decide whether the property was true or false of the category. In a control condition, each item contained a pronounceable nonword on top of the screen (e.g., “smalum”) and a single target letter on the bottom (e.g., “l”), and the participant had to decide whether the target letter was in the nonword. The fMRI contrast between experimental and control conditions showed an increase in activation in a region of interest (ROI) in the left fusiform gyrus. Importantly, this particular region is known to be involved in forming visual images of the objects (D’Esposito et al., 1997; Thompson-Schill, D’Esposito, Aguirre, & Farah, 1997), and may also be involved in the actual perceptual processing of objects. Activation in this same fusiform area has also been reported in other studies that require participants to verify object-property relations that are verbally stated (e.g., Chao et al., 1999; Kellenbach et al., 2001; Simmons et al., 2007). In addition, a whole-brain analysis by Kan et al. (2003) revealed only a single site of activation, one in the left fusiform gyrus close to the ROI. These results suggest that a region, or regions, in left fusiform gyrus—which are involved in visual-imagery and visual-object processing—contains the representations that are needed to answer verbal questions. In short, “knowing” a perceptual property amounts to re-perceiving it. This is perceptual embodiment.

While finding the left fusiform activation fits well with other published findings on property-verification, there are reasons to question the Kan et al. (2003) conclusion that the knowledge tapped in their study was specific to the visual modality. First, the failure of the whole-brain analysis to yield areas other than the left fusiform—say, in or near Wernicke’s area in left middle temporal cortex which is thought to be involved in amodal representations—may well be due to a lack of power since only seven participants were tested in the condition of interest (the study also included a second experimental condition with different participants)¹. Second, Kan et al. (2003) used a blocked design that averaged over True and False responses as well as over categories of living and non-living things. It is thus possible, say, that negative responses may rely on amodal representations or beliefs, (e.g., like the belief that “If property has not been found in an object representation by time t , it’s probably not true so respond negatively”) but the fMRI activations that could have revealed this abstract knowledge were obscured by the averaging process. Third, in retrospect, it would have been useful to have the left inferior frontal gyrus (IFG) serve as another ROI, as there is now evidence that this area is activated when verifying a property that has been rated as abstract (e.g., “lays eggs”) (Goldberg, Perfetti, Fiez, & Schneider, 2007)². Fourth, Kan et al. did not report any correlation between performance in the verification task and activation in the left fusiform gyrus. This leaves open the possibility that the area did not play a functional role verifying the properties, and that the necessary computations were being performed by other areas (whose activations were too small to detect because of a lack of power).

The current experiment uses the identical task and items as Kan et al. (2003)³ but: (1) tests a substantially larger number of participants; (2) uses an event-related design that allows us to separate the neural effects of type of response—True vs. False—and type of category—

¹Other studies supporting the embodiment hypothesis have routinely used more participants than Kan et al. (2003) and thus had more power, and often found active areas in addition to the fusiform (see Binder, Desai, Graves, & Conant, 2009). But these studies are still subject to some of the other criticisms mentioned.

²A study by Lee et al. (2002) did use ROIs that should reflect both modality-specific and amodal representations. However, this was a PET study that required the generation of properties rather than their verification. Some of the properties generated appear to have been quite complex (e.g., “lobsters are prized food, usually very expensive”), far more so than any of the properties employed in our study. So it is not clear how Lee et al.’s results bear on the present issues.

³We thank Irene Kan and Larry Barsalou for making these materials available to us.

living vs. non-living; (3) uses ROIs from the left fusiform gyrus and the left IFG, as markers of modality-specific and amodal representations/processes, respectively; and (4) reports correlations between activations and behavioral performance. To foreshadow our findings, we replicated the Kan et al. (2003) finding of left fusiform activation in verbal property-verification, but also found activation in the left IFG and regions in the middle temporal gyrus (in or near Wernicke's area). In addition, we report correlations between behavioral performance and activation in left IFG, an area associated with amodal representations and processes, but not with activation in the left fusiform. These results are more compatible with a theory that includes amodal components than with one that includes only modality-specific ones.

Methods

Participants

A total of 19 paid volunteers from the Columbia University community participated in the study. Participants were recruited and informed in accordance with IRB guidelines. One participant was excluded from further analysis due to excessive artifacts in the functional images, and 4 others were excluded because of excessive error rates, leaving a total of 14 participants (8 female, 6 male) with a mean age of 28.7 years (standard deviation: 10 years). (While the final number of participants is lower than desired, it is double the number used by Kan et al., (2003).) All participants were right-handed and had normal or corrected-to-normal vision, and had no history of neurological problems or traumatic brain injury.

Procedure

Participants performed the property verification task and a letter-verification task with nonwords as a control. In the property-verification task, the participant saw a fixation cross for 7 s at the beginning of each trial, followed by a category noun referring either to an animal, a plant, or an artifact (e.g., CAT). After 500 ms, while the category remained visible, a noun referring to a possible physical part appeared on the screen as well (e.g., WHISKERS). On 'false' trials, the property was semantically related to the concept, but not a physical part of it (e.g., PET). The participants were instructed to determine whether the property belonged to the category or not, and to respond as quickly as possible with a button press. From the onset of the property, participants had 1,500 ms to respond, after which the category and property were replaced by a fixation cross, indicating the beginning of the next trial. For both True and False trials in property-verification, half of them used categories of living things and half used non-living things. The control task was identical to the property-verification task except that the category name was replaced by a pronounceable nonword, the property by a single letter, and the task was to determine whether the single letter was contained in the nonword.

All participants began the experiment with the property-verification task and alternated between blocks of that task and blocks of the control task, completing four blocks of each for a total of 8 experimental blocks in total. Each block consisted of 50 trials (lasting approximately 8 min) of randomly intermixed True and False pairs. In total, participants saw 200 trials in each task. Each 8-min experimental block coincided with one run of MRI acquisition, so there were also 8 scanning runs in total.

Imaging

Data were acquired on a 1.5 Tesla GE Signa MRI scanner, in a slow event-related procedure. Stimuli were presented using E-Prime software (<http://www.pstnet.com/eprime.cfm>) and a back-projection screen that could be seen from the scanner bore via a mirror mounted to the head coil. Participants' responses were collected using an MR-

compatible button box. For each participant, we collected data in 8 scanning runs, each lasting approximately 8 min. Participants were allowed to rest for 30–60 s between runs.

Functional images were acquired using a standard echo-planar imaging sequence (26 contiguous slices with 4 mm thickness, in-plane resolution 64×64, field of view 224 mm, in-plane pixel size 3.5×3.5 mm, TR 2,000 ms, TE 40 ms, flip angle 90°). For the first 3 participants, 29 slices were acquired instead of 26, with all other scanning parameters remaining the same. After the experiment, participants remained in the scanner for two structural scans. A standard spoiled-gradient-recalled T1-weighted sequence with 1-mm³ resolution was used as the anatomical image. In addition, diffusion-weighted images were acquired, which were not part of the current study.

Data analysis

Data Preprocessing—Functional images were preprocessed and statistically analyzed using SPM 5 software (<http://fil.ion.ucl.co.uk>). They were corrected for differences in slice-time acquisition and corrected for motion, then registered to the anatomical scan. Functional and anatomical images were brought into standardized MNI space using a nonlinear normalization algorithm and interpolated to a new voxel size of 2×2×2 mm. The standardized functional images were smoothed with a three-dimensional Gaussian kernel (8 mm full-width at half-maximum).

Whole-brain Analysis of Task Effects—In a first-level event-related GLM analysis, we formed separate regressors of the stimulus-pair onset times for each condition and each of the 8 blocks. This resulted in 4 separate sets of property-verification regressors (1 set of regressors for each block) and 4 sets of control-task regressors. In each block of property-verification, we modeled the 4 task conditions (TRUE/FALSE X Living/Non-living) with separate regressors; each block of the control task was modeled with 2 regressors (TRUE/FALSE). Regardless of task, we also included one constant regressor for each of the 8 blocks. Because experimental blocks coincided with scanner runs, we used a different set of regressors for every block/run (even when the blocks contained the same task) in order to take into account variability in MRI signal strength across runs. Each regressor was convolved with a canonical approximation of the hemodynamic response function (a double gamma function provided in SPM). After estimating weights for each regressor using a maximum-likelihood approach (Friston et al., 1995), we computed basic contrast images comparing the mean of all property-verification regressors to the mean of all control task regressors, as well as more specific contrasts between the 4 conditions within the property verification task. Contrast images were then entered into a second-level analysis to determine areas of greater activation at the group level using a one-sample t-test. To control for possibly confounding effects of task difficulty, we entered mean reaction times and error rates for each task (property verification and control) as additional covariates at the second level, resulting in four additional regressors (one error rate and one reaction time regressor per task). This second analysis also allowed us to look for areas correlated with performance. Results were thresholded at $p < 0.01$ (FDR-corrected) and a minimum cluster size of 20 voxels. There are then two analyses, where the first one looks only at task effects—property-verification vs. control—and the second looks for specific effects within the property-verification task, e.g., living vs. nonliving categories.

Region-of-Interest Analysis—In addition to seeking effects at the level of the whole brain, we wanted to test our hypothesis regarding activation in regions of interest (ROIs) related to semantic processing and visual perception (see introduction). ROIs were constructed by placing spheres of 6mm radius around the maximum in the fusiform gyrus

reported by Kan et al. (2003) and 3 maxima reported by Badre et al. (2005), yielding a total of 4 ROIs.

The left fusiform region used by Kan et al. (2003) was originally reported by D'Esposito and colleagues (1997). We transformed the coordinates from the reported Talairach space ([-33, -48, -18]) into MNI space ([-35, -52, -19]) using the tal2icbm_spm transform (implemented in the Talairach daemon, available at <http://www.talairach.org/>). We chose the Badre et al. study because it contained a condition that was essentially property-verification, had a substantial sample size (22 participants), and showed separate activation sites for selection among semantic alternatives and for semantic representations themselves. The 3 selected ROIs were in the anterior left inferior frontal gyrus (BAs 44/47, at [-45, 27, -15] in MNI coordinates), a central region in the left inferior frontal gyrus (BA 45 at [-54, 21, 12]), and in the posterior middle temporal gyrus (BA 22 at [-51, -50, 5]).

Functional Connectivity Analysis—Finally, we determined whether functional connectivity between the above regions differed between property-verification and the control task. Toward this end, we extracted the timecourses from all voxels in the three ROIs taken from Badre et al. (2005), along with voxels in a 6-mm sphere surrounding the maximum in the left fusiform gyrus found in our initial whole-brain analysis. We then performed principal components analysis (PCA) on the timecourses of each region, resulting in a single timecourse for each region and session that was represented by the region's first eigenvector. PCA analysis extracts a representative signal that is robust to outlier voxels. Eigenvectors from the single sessions were then concatenated separately by task, resulting in one vector for each task in every region, i.e., 8 vectors for each participant (2 tasks x 4 ROIs). We then separately calculated partial correlation coefficients between each pair of the four ROIs, partialing out both the effects of the other two ROIs and of the global signal (extracted and averaged across the whole brain). We used partial correlations because we wanted to ensure that any correlation between a particular pair of ROIs could *not* be explained in terms of correlations between other pairs of ROIs. For each ROI pair, we then compared partial correlation coefficients between the property-verification and control tasks using paired *t*-tests, correcting for the six comparisons using the Bonferroni method.⁴

Results

Behavioral Results

We performed an analysis of variance (ANOVA) on reaction times (RTs) with the factors being task (property-verification vs. control) and response-type (True vs. False). There was a main effect of task ($F(1,13)=5.58, p<0.05$), RTs being faster in the control than the property-verification task (961 vs. 1019 ms). There was no main effect of response-type ($F<1$), but there was a significant interaction between task and response-type ($F(1,13)=16.93, p=0.01$), as False items took longer than True ones only in property-verification (994 vs. 1044 ms; $t(13)=2.1$). To test for category-specific effects, we performed an ANOVA just on the RTs in property-verification with the factors being category (living vs. non-living) and response (True vs. False). There was a significant main effect of category ($F(1,13)=30.79, p<0.01$), as RTs were faster to non-living than living categories (999 vs. 1038 ms), which held for both True and False responses. The results for error rates were similar to those for RTs. As determined by 2×2 ANOVAs, responses were more accurate in the control than the property-verification task (84.8% vs. 74.9%; $F(1,13)=10.34, p<0.01$), and more accurate with non-living than living categories (79.3%, vs. 71.2%; $F(1,13)=25.59, p<0.01$).

⁴The paired *t* tests were performed directly on the *r* values derived from the analysis. Such tests assume that the *r* values are normally distributed, which is a reasonable assumption with *r* values that fall below 0.5, as ours do.

fMRI Results

Whole-brain analysis—Compared with the control task, property-verification led to activation in a set of left-hemispheric temporal and frontal regions (see Table 1 and Figure 1, $p_{FDR} < 0.01$). In addition to finding activation in the left fusiform gyrus, as in Kan et al. (2003) and similar studies, task-related activity was seen in the: (1) left IFG (in the triangular aspect); (2) left posterior middle temporal gyrus (extending into the middle occipital gyrus), which overlaps with or borders on Wernicke's area; (3) left middle frontal gyrus; and (4) on the border of the left precuneus with the calcarine sulcus. The first two of these additional activations correspond to predicted areas that underlie amodal processes and perhaps amodal representations as well (but see discussion below). No regions were more activated during the control task than during property-verification ($p < 0.001$, uncorrected), and, at least at the whole-brain level, no regions correlated significantly with reaction time or accuracy ($p < 0.001$, uncorrected, but see Relation to Accuracy, below).

ROI analyses—We conducted two different ROI analyses. The first centered on the left fusiform region used by Kan et al. (2003) and compared the mean activation in all property verification conditions against all control task conditions. This analysis revealed a marginally significant activation ($t(12) = 1.74$, $p = 0.054$, one-tailed). Thus, we find support for the Kan et al. (2003) hypothesis that perceptual knowledge is represented in areas thought to be modality-specific. The second set of 3 ROIs was composed of regions in the anterior and central left inferior frontal gyrus and the left middle temporal gyrus (Badre et al., 2005). We found that both the anterior and central IFG ROIs showed more activation in our property-verification than our control task ($t(12) = 4.43$, $p < 0.005$, in the anterior ROI; $t(12) = 4.72$, $p < 0.001$, in the central IFG ROI). Activation in the posterior middle temporal gyrus, though nominally weaker, was still significantly greater in property-verification, than control ($t(12) = 2.74$, $p < 0.05$). Since the second set of ROIs have previously been associated with semantic but not perceptual processing (i.e. in Badre et al., 2005), their activation during property verification suggests that part of the knowledge required in this task may have been amodal. Even though it is problematic to infer cognitive functions from activation in a specific brain region without further quantitative analysis (as in the approach formulated in Poldrack, 2006), the perceptual controls in Badre et al. make it unlikely that these regions are more involved in perceptual than amodal processing.

Effects of Category—We tested for category-specific effects by comparing trials with a living category to trials with a nonliving category within areas activated by property-verification (as defined by a mask created from the contrast property-verification minus control, thresholded liberally at $p < 0.001$ and uncorrected for multiple comparisons--this is the second level analysis described earlier). The contrast living > nonliving showed greater activation for the living category in the left fusiform gyrus and extended into the parahippocampal gyrus (BAs 36/37, -30, -40, -14, $p < 0.01$). We also found increased activation for the living category in the left precuneus, extending into the posterior cingulate and calcarine sulcus (BAs 30/17, -12, -52, 8, $p < 0.05$). Nonliving categories did not activate any areas more strongly than living categories, even when searching in all brain voxels (at a liberal threshold of $p < 0.001$, uncorrected).

Since living category items led to longer RTs than nonliving category items, it is possible that these activation effects could be due to increased time-on-task. An alternative is that living things contain more perceptual properties than nonliving things, hence activate perceptual areas more. The latter hypothesis has figured centrally in the neuropsychology literature, where it has been offered as an explanation of why many patients with agnosia are impaired in recognizing and categorizing living things, but not nonliving ones (e.g., Warrington & Shallice 1984; Farah & McClelland, 1991). However, the most up-to-date

systematic review of this literature offers no support at all for this hypothesis (Capitani, Laiacona, Mahon & Caramazza, 2003).

In contrast to the category effects, we found no significant differences between True and False trials in the property verification task (at the whole brain level with a threshold of $p < 0.001$, uncorrected).

Relation to Accuracy—In addition to the strong task-related activation in the anterior IFG, we also found a significant correlation between activation in this area (relative to the control task) and behavioural performance. Specifically, we did a multiple regression analysis of our performance measures – RT and error-rate differences between property-verification and control – on the IFG activation difference between property-verification and control. There was a significant difference in regression weights for error rates (r-equivalent $t(12) = 2.35$, $p < 0.05$). This indicates that, relative to the control task, participants with higher accuracy in property-verification showed more activation in the left anterior IFG, strengthening the claim that this amodal area played a functional role in property verification. In contrast, there was no significant correlation between performance and activation in the left fusiform gyrus⁵. Using ROIs and correcting for baseline differences in performance and activation by subtracting the control task values seems to have been necessary to find this effect (see whole-brain results above). We did not find any significant correlation between brain activity and RTs. However, the obtained RTs may not have been very sensitive, since responses longer than the 1,500 ms response window were not recorded.

Functional connectivity—Compared to the control task, during property-verification the partial correlation coefficient between activity in anterior IFG and the fusiform gyrus increased significantly ($t(12) = 3.26$, $p < 0.05$ (in property-verification, $r = 0.186 \pm 0.028$; in the control task, $r = 0.121 \pm 0.032$). Thus there is evidence for a functional connection between left-hemisphere IFG and fusiform gyrus. There were no other significant partial correlations between activations in different areas (all $ps > 0.1$).

Discussion

Summary

Our methodological changes of the Kan et al (2003) study led to results that now support a role for amodal representations or processes in property-verification. First, increasing the power of the experiment changed the results of a whole-brain analysis such that now we obtained activations in areas that are associated with amodal representations or processes, particularly areas in left IFG, and in or near Wernicke's area. This result is in keeping with numerous other recent neuroimaging studies of semantic memory (see the recent meta-analysis of Binder et al., 2009). Second, using ROIs presumed to reflect amodal processing, and finding significant activation in them, strengthens the case for an amodal component in the verification of perceptual properties. Third, the correlations between activations in the anterior IFG and accuracy in the property-verification task suggest that this area played a functional role in performing the task; this is in contrast to the lack of a correlation between activations in the fusiform gyrus and performance. In addition, we found that the IFG was functionally connected to the fusiform gyrus.

⁵Ebisch et al. (2007) also reported a correlation between activation in left IFG and performance in a semantic judgment task, though their judgments appear to be more difficult than those required in the present study.

An amodal hypothesis for “simple” property verification

Given our evidence for the involvement of amodal representations or processes, it is worth reconsidering the *functional* role of the left fusiform gyrus in property-verification studies like Kan et al. (2003) and the present one. Consider the hypothesis that while the fusiform-gyrus is indeed a perceptual area it played little role in our experiment. Perhaps the amodal representations and processes in or near Wernicke’s area and IFG are sufficient for answering simple questions about properties—where “simple” implies that no imagery is required (Kosslyn, 1976); but these representations may be *closely linked* to modality-specific ones in the fusiform gyrus that are needed when property verification requires visual imagery (e.g., “parrot-curved beak”). In our task, activation of the relevant amodal representation might have led to concurrent activation of the associated modality-specific representation, though the latter would have played *no functional role in our study*. This hypothesis is similar to the recent proposals of Mahon and Caramazza (2008, 2009), who suggest that activations in motor areas may not reflect the computations involved when viewing manipulable objects, but instead reflect a spread of activation from the amodal areas that do the real neural work. By analogy, in our study activation from amodal areas could spread to perceptual processing regions such as the fusiform gyrus.

We are not proposing that the left fusiform *never* has a functional role in property verification. Rather, that area may help mediate visual imagery, which may be needed to verify complex perceptual properties as in “parrot-curved beak?” (D’Esposito et al., 1997, Thompson-Schill et al., 2003). And when modality-specific representations are playing a functional role—which may be accompanied by a greater level of activation in the fusiform gyrus than that obtained in studies like the present one--the participant may experience visual images. Importantly, there is little mention of participants reporting visual images in property-verification studies like the present one. While modality-specific accounts need not require visual imagery, the absence of imagery here underscores that the perceptual demand of our task was simple enough to be performed by an amodal system.

So far, we have assumed that the activations we found in the IFG and MTG reflect abstract and modality-independent processing. However, recent imaging studies have tied the latter area (the middle temporal gyrus) to the processing of biological and mechanical motion (e.g., Wheatley et al., 2007). This study and others (Lin et al., 2011) have found dissociable regions around the middle and superior temporal gyri relating to the (presumably perceptual) representation of different kinds of motion. The peak we found was more posterior and superior (bordering on the middle occipital gyrus) than is typically reported in studies of amodal semantic processing and may therefore indicate that the activation was in fact related to motion processing. This would mean that, along with the fusiform gyrus, we found activation in two modality-specific areas, rather than just one. However, in contrast with previous findings, the activation we found in the middle temporal gyrus was independent of noun category (since there was no activation difference in this region between animate and inanimate nouns). Further, we found significant activation in an ROI of the left MTG, in a region both closer to Wernicke’s area and already shown to activate in another semantic task (i.e., Badre et al., 2005). Therefore, we suggest that the activation found in our study does not simply reflect activation of motion properties of the presented nouns, but rather that it also represents amodal information.

Another kind of amodal hypothesis arises if we assume that the representations in the fusiform gyrus did play a functional role, but those representations were in fact amodal, i.e., the fusiform gyrus supports amodal representations as well as modality-specific ones. This assumption fits with the neuropsychological findings of Büchel, Price & Friston (1998), who tested blind participants with damage in the ventral temporal cortex and normally-sighted people in a reading task; their results were most compatible with the hypothesis that the

ventral temporal cortex is primarily an association area that integrates different types of sensory information. (See also Mahon et al. (2009) who found that ventral temporal cortex responds to category-level semantic information in congenitally blind people.) This amodal hypothesis, however, would lead one to expect correlations between activations in the fusiform gyrus and behavioural performance, which were not obtained. Indeed, we have used this lack of correlation as evidence against a modality-specific hypothesis of property verification.

Problems with a purely amodal account

As just noted, our amodal hypothesis is partly based on the lack of a correlation between activations in the fusiform gyrus and behavioural performance, in the face of a significant correlation between activations in IFC and behavioural accuracy. This kind of negative evidence has the usual liabilities, particularly that it may reflect a lack of statistical power; though our number of participants was greater than that of Kan et al. (2003) it was still only a modest 13. Perhaps even more important, we have made much of the activations in IFC in arguing against a purely modality-specific account, but these activations could reflect primarily general retrieval processes, rather than semantic representations or semantic-specific retrieval processes. Thus one of our ROIs has often been argued to mediate the function of selecting among retrieved representations regardless of their type (e.g., Thompson-Schill et al., 2003). If IFC mediated only general retrieval operations, the representations themselves must be elsewhere. Thus it is still possible that the representations involved were modality-specific ones in the fusiform gyrus, but that they were retrieved by top-down activation from domain-general processes. Even in this case, a theory of property-verification must include processes that are not tied to a specific modality.

Finally, it is worth noting that the most successful computational models of brain-damage and semantic deficits assume that the areas affected involve amodal or multi-modal representations, not modality-specific ones (e.g., Rogers & McClelland, 2004). Such computational evidence combined with linguistic analyses cited in the Introduction (Miller & Johnson-Laird, 1971) and the experimental findings presented here, might give researchers pause about embracing the embodiment theory of semantic memory.

Acknowledgments

This research was supported in part by NIH grant AG015116-12.

We thank Murray Grossman for his helpful comments on an earlier version of this manuscript.

References

- Allport, DA. Distributed memory, modular subsystems and dysphasia. In: Newman, SK.; Epstein, R., editors. *Current Perspectives in Dysphasia*. New York: Churchill Livingstone; 1985. p. 207-44.
- Badre D, Poldrack RA, Paré-Blagoev EJ, Insler RZ, Wagner AD. Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*. 2005; 47:907–918. [PubMed: 16157284]
- Barsalou LW, Simmons WK, Barbey AK, Wilson CD. Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*. 2003; 7:84–91. [PubMed: 12584027]
- Beauchamp MS, Haxby JV, Jennings JE, DeYoe EA. An fMRI version of the Farnsworth-Munsell 100-Hue test reveals multiple color-selective areas in human ventral occipitotemporal cortex. *Cerebral Cortex*. 1999; 9:257–263. [PubMed: 10355906]
- Binder JR, Desai RH, Graves WW, Conant LL. Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*. 2009; 19:2767–2796. [PubMed: 19329570]

- Büchel C, Price C, Friston K. A multimodal language region in the ventral visual pathway. *Nature*. 1998; 394:274–277. [PubMed: 9685156]
- Capitani E, Laiacona M, Mahon B, Caramazza A. What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology*. 2003; 20(3/4/5/6): 213–261. [PubMed: 20957571]
- Chao L, Haxby JV, Martin A. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*. 1999; 2:913–919.
- Collins AM, Loftus EF. A spreading activation theory of semantic processing. *Psychological Review*. 1975; 82:407–428.
- D'Esposito M, Detre JA, Aguirre GK, Stallcup M, Alsop DC, Tippet LJ, Farah MJ. A functional MRI study of mental image generation. *Neuropsychologia*. 1997; 35:725–730. [PubMed: 9153035]
- Dreyfuss, M.; Smith, EE.; Cook, P.; McMillian, CT.; Bonnet, MF.; Richmond, L.; Grossman, M. Neural networks underlying object knowledge. (submitted)
- Ebisch SJH, Babiloni C, Del Gratta C, Ferretti A, Perrucci MG, Caulo M, Romani GL. Human neural systems for conceptual knowledge of proper object use: A functional magnetic resonance imaging study. *Cerebral Cortex*. 2007; 17:2744–2751. [PubMed: 17283202]
- Hauk O, Johnsrude I, Pulvermüller F. Somatotopic representation of action words in human motor and premotor cortex. *Neuron*. 2004; 41:301–307. [PubMed: 14741110]
- Goldberg RF, Perfetti CA, Fiez JA, Schneider W. Selective retrieval of abstract semantic knowledge in left prefrontal cortex. *The Journal of Neuroscience*. 2007; 27:3790–3798. [PubMed: 17409243]
- Goldberg RF, Perfetti CA, Schneider W. Perceptual knowledge retrieval activates sensory brain regions. *J Neurosci*. 2006; 26:4917–4921. [PubMed: 16672666]
- Kan IP, Barsalou LW, Solomon KO, Minor JK, Thompson-Schill SL. Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge. *Cognitive Neuropsychology*. 2003; 20:525–40. [PubMed: 20957583]
- Kellenbach ML, Brett M, Patterson K. Large, colorful, or noisy? Attribute-and modality-specific activations during retrieval of perceptual attribute knowledge. *Cognitive, Affective and Behavioral Neuroscience*. 2001; 1:207–221.
- Kosslyn SM. Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval times. *Memory and Cognition*. 1976; 4:291–297.
- Lee AC, Graham KS, Simons JS, Hodges JR, Owen AM, Patterson K. Regional brain activations differ for semantic features but not categories. *NeuroReport*. 2002; 13:1497–1501. [PubMed: 12218693]
- Lin N, Lu X, Fang F, Han Z, Bi Y. Is the semantic category effect in the lateral temporal cortex due to motion property differences? *NeuroImage*. 2011; 55:1853–1864. [PubMed: 21256232]
- Mahon BZ, Caramazza A. A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology: Paris*. 2008; 102:59–70.
- Mahon BZ, Anzellotti S, Schwarzbach J, Zampini M, Caramazza A. Category-specific organization in the human brain does not require visual experience. *Neuron*. 2009; 63:397–405. [PubMed: 19679078]
- Mahon BZ, Caramazza A. Concepts and categories: A cognitive neuropsychological perspective. *Annual Review of Psychology*. 2009; 60:1–25.
- Martin A, Haxby JV, Lalonde FM, Wiggs CL, Ungerleider LG. Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*. 1995; 270:102–105. [PubMed: 7569934]
- Miller, GA.; Johnson-Laird, PN. *Language and perception*. 1. Cambridge, MA: Harvard University Press; 1976.
- Mummery CJ, Patterson K, Hodges JR, Price CJ. Functional neuroanatomy of the semantic system: Divisible by what? *Journal of Cognitive Neuroscience*. 1998; 10:766–777. [PubMed: 9831743]
- Poldrack R. Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Science*. 2006; 10(2):59–63.
- Rogers, TT.; McClelland, JL. *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge, MA: MIT Press; 2004.

- Simmons WK, Ramjee V, Beauchamp MS, McRae K, Martin A, Barsalou LW. A common neural substrate for perceiving and knowing about color. *Neuropsychologia*. 2007; 45:2802–2810. [PubMed: 17575989]
- Simmons WK, Ramjee V, McRae K, Martin A, Barsalou LW. fMRI evidence for an overlap in the neural bases of color perception and color knowledge. *Neuroimage*. 2006; 31:S182.
- Smith, EE.; Medin, DL. *Categories and concepts*. Cambridge, MA: Harvard University Press; 1981.
- Smith EE, Shoben EJ, Rips LJ. Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*. 1974; 81:214–241.
- Thompson-Schill SL. Neuroimaging studies of semantic memory: Inferring “how” from “where”. *Neuropsychologia*. 2003; 41:280–292. [PubMed: 12457754]
- Thompson-Schill SL, D’Esposito M, Aguirre GK, Farah MJ. Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proceedings of the National Academy of Sciences*. 1997; 94:14792–14797.
- Wheatley T, Milleville SC, Martin A. Understanding Animate Agents: Different Roles for the Social Network and Mirror System. *Psychological Science*. 2007; 18(6):469–474. [PubMed: 17576256]

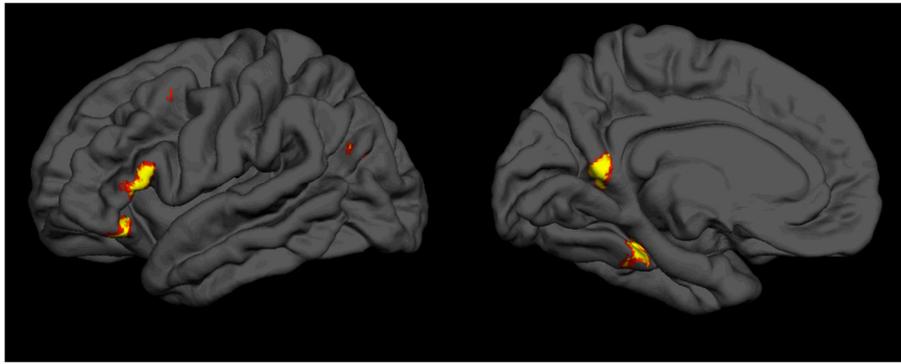


Figure 1. Activations related to property verification, rendered on a standard participant's brain in MNI space. Medio-ventral view (right) shows a maximum in the left fusiform gyrus. The lateral view (left) shows activation maxima in the left posterior middle temporal and left inferior and middle frontal gyri. Results are thresholded at an alpha-level of 0.05, FDR-corrected for multiple comparisons, and exclude clusters containing fewer than 15 voxels.

Table 1

Regions activated in property verification compared to the control condition.

Property Verification > Control Task									
k	x	y	z	pFDR	Z	BA	H	region	
169	-32	-36	-18	.003	4.89	37	L	fusiform gyrus (FG)	
123	-10	-56	6	.005	4.23	30	L	precuneus	
	-18	-52	8	.007	3.84	30	L	calcarine sulcus	
24	-38	-70	26	.005	4.21	39	L	middle occipital gyrus	
75	-38	32	-16	.006	4.16	47	L	IFG (orbital)	
29	-44	14	44	.007	3.85	9	L	middle frontal gyrus	
45	-50	28	4	.008	3.65	45	L	IFG (triangular)	
54	-54	-64	14	.007	4.01	22	L	middle temporal gyrus	

Note: k: cluster size in voxels (2mm^3 voxel size). Z: equivalent Z-score. BA: Brodmann's area. H: hemisphere. Coordinates are in MNI space. pFDR: false-discovery-rate-corrected p-value.