

ROLE OF MENTAL IMAGERY IN A PROPERTY VERIFICATION TASK: fMRI EVIDENCE FOR PERCEPTUAL REPRESENTATIONS OF CONCEPTUAL KNOWLEDGE

Irene P. Kan

University of Pennsylvania, Philadelphia, USA

Lawrence W. Barsalou

Emory University, Atlanta, GA, USA

Karen Olseth Solomon

Willamette University, Salem, MA, USA

Jeris K. Minor and Sharon L. Thompson-Schill

University of Pennsylvania, Philadelphia, USA

Is our knowledge about the appearance of objects more closely related to verbal thought or to perception? In a behavioural study using a property verification task, Kosslyn (1976) reported that there are both amodal and perceptual representations of concepts, but that amodal representations may be more easily accessed. However, Solomon (1997) argued that due to the nature of Kosslyn's stimuli, subjects may be able to bypass semantics entirely and perform this task using differences in the strength of association between words in true trials (e.g., cat-whiskers) and those in false trials (e.g., mouse-stinger). Solomon found no evidence for amodal representations when the task materials were altered to include *associated* false trials (e.g., cat-litter), which require semantic processing, as opposed to associative strategies. In the current study, we used fMRI to examine the response of regions of visual association cortex while subjects performed a property verification task with either associated or unassociated false trials. We found reliable activity across subjects within the left fusiform gyrus when true trials were intermixed with *associated* false trials but not when true trials were intermixed with *unassociated* false trials. Our data support the idea that conceptual knowledge is organised visually and that it is grounded in the perceptual system.

One of the leading theories of the organisation of semantic knowledge holds that conceptual information is distributed across distinct attribute domains, such as vision, touch, and action (Allport,

1985). Numerous investigators have drawn distinctions between visual knowledge and nonvisual (e.g., functional) knowledge, in areas ranging from language acquisition (e.g., Gentner, 1978; Nelson,

Requests for reprints should be addressed to Irene P. Kan, Department of Psychology, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA 19104-6196, USA (Email: ikan@psych.upenn.edu).

Supported by grants from National Institute of Health (NIH R01 MH60414) and Searle Scholars Program awarded to STS, and grant from National Science Foundation (SBR-9905024) awarded to LWB. We thank Jessica Pickard for assistance with data collection, Geoff Aguirre and Joseph Kable for assistance with data analysis, and two anonymous reviewers for helpful comments on an earlier version of this manuscript.

1974) to object categorisation (e.g., Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Many neuroimaging studies have supported the distinction between visual and nonvisual knowledge (see Thompson-Schill, 2003, for a review). One interpretation of these studies is that our knowledge about the visual attributes of an object is represented differently from our knowledge of other, nonvisual attributes. This claim is related to a fundamental question that has been debated by philosophers (e.g., Locke, 1690/1959) and psychologists (e.g., Kosslyn, 1994; Pylyshyn, 1981) for centuries: To what extent does conceptual knowledge rely on perceptual representations, as opposed to propositional, amodal representation? That is, are the underlying representations descriptive (i.e., more akin to verbal thoughts and nonperceptual) or are they depictive (i.e., perceptual and possess visual-spatial qualities)? Until the mid-twentieth century, the view that conceptual knowledge is perceptually based was widely accepted (e.g., Locke, 1690/1959; Price, 1953). However, this assumption was called into question, and the nature of the internal representations of conceptual knowledge has undergone much scrutiny over the past few decades.

One approach that researchers have taken to address these questions is to examine the processes involved in mental imagery. It is assumed that in order to create a mental image of an object, access to conceptual knowledge is necessary. That is, in order to create a mental image of an apple, one must retrieve the stored knowledge about the concept "apple." Thus, it seems that understanding the nature of mental imagery may help us address more fundamental questions about the representation of long-term conceptual knowledge, especially knowledge about visual attributes. According to depictive theories of conceptual knowledge, mental imagery involves the same representations and mechanisms normally used during visual perception (e.g., Kosslyn, 1994; Shepard & Cooper, 1982; Zatorre, Halpern, Perry, Meyer, & Evans, 1996). That is, conceptual representations are grounded in the perceptual systems and that information is represented in a spatial and pictorial format. One line of evidence in support of this theory came from

visual scanning experiments. In Kosslyn, Ball, and Reiser's (1978) classic study, subjects were instructed to focus their attention on one part of a mental image and "move" that part to the location of some other part of the image. They reported that the time taken to complete a mental scan was directly proportional to the physical distance between the two locations. Kosslyn and colleagues inferred from these findings that mental imagery and visual perception must share similar mechanisms.

However, the validity of these conclusions has been questioned. For example, Pylyshyn (1981) argued that subjects interpreted the task demands as simulating the use of visual-spatial knowledge. Similarly, Intons-Peterson (1983) suggested that subjects were responding to experimenters' expectations, and that their performance was biased by the experimenters' presumptions. These investigators have proposed that subjects' expectations resulted in performance that appeared to support a perceptual-based theory, even though the underlying representations were not perceptual. Instead, Pylyshyn and others have argued that there is a complete independence between the perceptual system and the conceptual system, such that when an object is perceived, its visual features will be represented in the perceptual system, and the conceptual system will reinterpret that information as a list of propositional features. For example, on perceiving an apple, the perceptual system represents it in a visual form, and the conceptual system reinterprets that information into an amodal feature list, with entries such as "shiny," "red," "round." Proponents of propositional theories argued that since much of cognition depends on a "language of thought," it is only natural that mental images are represented propositionally, and that this amodal list is retrieved when conceptual knowledge is accessed (Fodor & Pylyshyn, 1988; Pylyshyn, 1984). However, as Barsalou (1999) noted, nearly all evidence for amodal theories is indirect, relying either on theoretical arguments, or on the ability to implement amodal representations formally and computationally.

The only exception seems to be Kosslyn's (1975, 1976) property verification experiments, in which

he investigated whether visual imagery and other forms of internal representations can be differentiated. In this task, subjects were asked to determine whether certain properties are true or false of certain concepts. For example, when presented with the concept word "cat," subjects were asked to verify whether the property "claws" is true of that concept. To determine whether individuals would consult amodal knowledge or perceptual knowledge when verifying properties, Kosslyn (1976) manipulated task instructions within the same group of subjects. For the first block of trials, subjects did not receive explicit instructions on how to perform the task. They were simply told to verify the properties as quickly as possible without sacrificing accuracy. For the second block of trials, subjects were given explicit imagery instructions. They were told to first form a vivid mental image and then to verify each property by consulting their mental picture. Furthermore, Kosslyn systematically varied two other factors: size of a property relative to its concept and the associative strength between a property and its concept.

Under neutral instructions, subjects were significantly faster at verifying properties than under the imagery condition. Furthermore, reaction times collected under the neutral instructions were correlated with associative strength between properties and their concepts, which support theories of amodal representations. On the other hand, reaction times collected under the imagery condition were not correlated with associative strength; instead, the reaction times were inversely correlated with property size. These results support theories of perceptual representations. In sum, Kosslyn concluded that there are both amodal and perceptual representations of concepts, but that amodal representations may be more easily accessed.

Although these findings have been cited extensively, recent work challenges the extent to which these findings can elucidate the representation of conceptual information. Solomon (1997) questioned the nature of the stimuli used in Kosslyn's (1976) original studies (see also Barsalou, Solomon, & Wu, 1999), noting that the true trials and the false trials actually differed on two dimensions: (1) whether the properties belong to the concepts,

and (2) whether the properties are associated with the concepts. One way of assessing associative strength between two concepts is how frequently the two concepts co-occur in the same context. The more frequently they occur together, the more associated they are. Another way of thinking about associative strength is how readily a concept becomes available when another concept is accessed. For example, consider the two concept-property pairs "cat-head" and "mouse-stinger." The words in the first pair are associated because these words often appear together within the same context. On the other hand, the words "mouse" and "stinger" rarely occur within the same context, and it is unlikely that these two words are associated with each other. In other words, not only do these word pairs differ on whether the property belongs with the concept, they also differ on associative strength.

The inadvertent variation in association strength between false and true trials may have had consequences for the results reported by Kosslyn (1976). Solomon (1997) pointed out that while all of the true properties in Kosslyn's original study were associated to their concepts (e.g., cat-head), all of the false properties were unassociated to their concepts (e.g., mouse-stinger). It was possible for subjects in the neutral condition to perform the task based on a simple word association strategy. That is, subjects might have adopted a strategy of answering "yes" if the concept and property are at all associated and answering "no" if the concept and property are unassociated. Furthermore, it has been shown that when using a simple word association strategy, access to conceptual knowledge is not obligatory (Glaser, 1992). Glaser proposed a model that includes a direct connection between the printed words and the verbal lexicon, and this route is independent from the conceptual system. If subjects in Kosslyn's experiment were in fact not accessing conceptual knowledge to perform the task in the neutral condition, then the findings from that experiment do not bear on the question of the nature of conceptual knowledge representations.

Solomon (1997) replicated Kosslyn's property verification experiments (1975, 1976), with one

important variation: To ensure that conceptual processing would be engaged in all property verification trials, Solomon replaced the unassociated false trials with *associated* false trials (e.g., steak-grill) in one condition of the experiment. Solomon reasoned that subjects would not be able to rely on a simple word-association strategy when the concepts and properties of both true and false trials were equivalently associated. Under these conditions, subjects' performance was predicted by perceptual factors (e.g., size) and not linguistic factors (e.g., associative strength). Contrary to Kosslyn's (1976) results, Solomon obtained this pattern of data even when subjects received neutral instructions (i.e., associated-neutral subjects). The author interpreted these data as support for the idea that conceptual knowledge is grounded in perceptual simulation. The present experiment follows from this study, but additionally addresses differences in neural activity between associated and unassociated conditions.

One approach to investigating the relation between conceptual knowledge and perception has been to compare neural mechanisms of each. This approach has been motivated, in part, by the unresolved questions from behavioural experiments. Neuropsychological studies of patients with perceptual and imagery deficits as a result of brain damage have been used to bolster the hypothesis that mental imagery and perception share common processes (e.g., Bisiach & Luzzatti, 1978; Farah, Soso, & Dasheiff, 1992; Servos & Goodale, 1995). Recently, functional neuroimaging has been used to investigate this question in normal subjects. Some studies have reported activation in visual cortex during mental imagery (e.g., Charlot, Tzourio, Zilbovicius, Mazoyer, & Denis, 1992; D'Esposito et al., 1997; Goldenberg, Podreka, Steiner, Franzen, & Deecke, 1991; Goldenberg, Podreka, Steiner, & Willmes, 1987; Roland & Friberg, 1985), in extrastriate and occasionally primary visual areas (see Kosslyn & Thompson, 2000, for a full review of this topic). As this approach will be used in the present investigation, we briefly review some relevant findings here.

Using functional magnetic resonance imaging (fMRI), D'Esposito and colleagues (1997) used a

mental image generation task to examine the neural bases of mental imagery. Subjects listened to alternating blocks of concrete words (e.g., apple, horse) and abstract words (e.g., treaty, guilt) while in the scanner. Subjects were instructed to generate mental images of the concrete words' referents, and they were told to listen passively to the abstract words. D'Esposito and colleagues reported that in neurologically intact adults, visual association cortex was recruited during mental imagery. Specifically, when compared to passive listening of abstract words, generating images in response to concrete words resulted in significantly greater activity that was asymmetrically lateralised to the left fusiform gyrus, left premotor area, and the left anterior cingulate gyrus. These data were also consistent with other neuroimaging studies of mental imagery (e.g., Goldenberg et al., 1987, 1991).

Thompson-Schill, Aguirre, D'Esposito, and Farah (1999) reported similar findings in their investigation of the neural bases of semantic knowledge using fMRI. Subjects in their experiment were asked a series of yes/no questions about visual or nonvisual characteristics of living and nonliving objects. For example, the following questions were posed to the subjects: visual/living—"Does a parrot have a curved beak?"; nonvisual/living—"Are pandas found in China?"; visual/nonliving—"Are bows of violins longer than violins?"; and non-visual/nonliving—"Does a toaster use more electricity than a radio?" Consistent with D'Esposito et al.'s (1997) results, Thompson-Schill and colleagues also found differential patterns of activations between visual knowledge retrieval and nonvisual knowledge retrieval. Specifically, greater left fusiform gyrus activity was found when subjects engaged in visual knowledge retrieval than when they executed nonvisual knowledge retrieval. Additionally, in one subject who participated in both experiments, when activation during visual knowledge retrieval was compared to the explicit imagery task used by D'Esposito et al., considerable overlap in the left fusiform gyrus was found between these two tasks. On the basis of these results, the left fusiform gyrus was a crucial region of interest in the present investigation.

The aim of the present study was to demonstrate the involvement of visual cortical regions during semantic processing—even with no explicit imagery instructions—and furthermore to show that such involvement depends on the extent to which semantic knowledge, as opposed to lexical association information, is required to perform the task. Specifically, we used fMRI to examine the involvement of the left fusiform gyrus in the property verification task. Activity in the visual association cortex during property verification would support the assertion that semantic knowledge is indeed grounded in the perceptual system. However, this prediction only applies to tasks that require the retrieval of semantic knowledge. Thus, we hypothesised that different patterns of activation would be observed depending on whether the false trials in the experiment were associated or unassociated pairs. To test this hypothesis, subjects were randomly assigned to the associated condition, in which both true and false trials were associated pairs (as in Solomon, 1997), or the unassociated condition, in which true trials were associated but false trials were unassociated pairs (as in Kosslyn, 1976). Following from Solomon, we predicted that left fusiform activity during property verification would only be observed in the associated condition because this condition forces subjects to retrieve semantic knowledge. In contrast, activation of the left fusiform gyrus during property verification would not be expected to reliably occur in the unassociated condition, in which subjects are able to adopt a word association strategy to perform the task. If semantic knowledge is represented amodally (Pylyshyn, 1981, 1984), then no activation in the visual association cortex is expected in either condition.

METHODS

Subjects

Fourteen subjects from the University of Pennsylvania participated in this experiment. Subjects were randomly assigned to either the associated or unassociated task condition. Subjects in the associated

group were two males and five females, aged 19–21 years (mean age = 19.9 years). Subjects in the unassociated group were five males and two females, aged 18–22 years (mean age = 20.6 years). All subjects met the following inclusion criteria: They were (1) high-school educated, (2) native English speakers, and (3) right-handed. General exclusionary criteria were (1) history of neurological or psychiatric illness or (2) current use of medication affecting the central nervous system (e.g., psychotropic drugs). Informed consent was obtained from all subjects. Each subject was paid \$20 for his or her participation.

Materials

Property verification task

Each subject performed a property verification task, in which the task was to determine whether a property (e.g., frosting) was part of a concept (e.g., cake). Subjects were told that each stimulus pair was arranged in a top–bottom configuration and that all concept words were presented on the top and all property words were presented on the bottom. All of the concept–property pairs used in this experiment were originally developed by Solomon (1997). To construct the 100 true trials, 100 properties (e.g., hood, udder, drawer) were paired with 100 concepts (e.g., car, cow, dresser) from 18 different superordinate categories (e.g., vehicle, mammal, furniture), and all properties were limited to physical parts of concepts. For example, “hill” could be loosely considered a property of “ant,” but here it was considered a false property because it was not a physical part of the concept. These pairings varied widely along both association strength and size of property, and no property or concept was ever repeated.

For the false trials, two sets of materials were constructed. The associated set was constructed by pairing 100 concepts with 100 associated false properties (i.e., canary–sing). The unassociated set was constructed by pairing the same 100 concepts with a random order of the same 100 false properties (i.e., canary–wine), thereby creating unassociated pairs, but also keeping the concepts and

properties constant. No false property or concept was repeated or used as a true property or concept.

To assess the association strength of both the true and false pairs, an independent group of 20 subjects from University of Chicago rated the association strength of all possible pairs. For this task, subjects first read a concept word, then a property word, and then rated how much the property word had come to mind while reading the concept word. Subjects made their rating from a scale of 1 (*did not come to mind*) to 9 (*immediately came to mind*). The true pairs varied widely in association strength, with the 50 most-associated properties averaging 7.5 and the 50 least-associated properties averaging 4.6. The two kinds of false pairs differed markedly, with the associated false pairs averaging 5.0 and the unassociated false pairs averaging 1.2.

True pairs were also chosen to represent a wide variety of sizes of properties of the concepts. For example, a cat's head is larger than a cat's claw and therefore a head should be easier to resolve on an image of a cat than a claw. To assess the size of the property on the concept, 12 new subjects imagined each concept and estimated the percentage of the concept that the property occupied. For example, subjects were asked, "What percentage of the total volume of a fox is its nose?" Subjects gave a rating from 0 to 100 for each concept–property pair. The largest 50 properties averaged 29.7%, whereas the smallest 50 properties averaged 10.6%. Concept–property pairs were also chosen so that association strength and the size of the property was not correlated ($r = .13, p > .25$).

To summarise, there were three types of property verification trials—true, associated false, and unassociated false. The correct response for a true trial was "yes," and the correct answer for a false trial was "no." Subjects in both the associated and the unassociated groups encountered the same true trials (e.g., goat–ear). The critical trials that differed between the two groups were the false trials. Subjects in the associated condition received associated false trials (e.g., stapler–paper), and subjects in the unassociated condition received unassociated false trials (e.g., stapler–vegetable). Although both associated false trials and unassociated false trials required a "no" response, the relationships between

the stimulus words differed in the two conditions. In the associated false trials, even though the two words were associated with each other, the property was not part of the object, but in the unassociated false trials, the property was neither part of the object nor an associate of the concept.

Baseline letter verification task

In the letter verification task, subjects determined whether a single letter (e.g., "m") was part of a pronounceable nonword letter string (e.g., "smalum"). As in the property verification task, the stimulus pair was arranged in a top–bottom configuration, where the pronounceable nonword appeared on the top and the single letter appeared on the bottom. Average number of letters in the nonwords was matched to those of the experimental trials.

Procedures

The procedure used in the associated and unassociated groups was identical, with the only difference being the stimuli in the critical false trials, as described above. Each trial began with a fixation point presented in the middle of the screen for 200 ms. The concept word (or nonword) followed the fixation and was presented for 2800 ms and remained on the screen for the duration of the trial. The stimulus onset asynchrony between the concept word (or nonword) and the property word (or letter) was 1300 ms. In other words, the total exposure time for the concept word (or nonword) was 2800 ms, and the total exposure time for the property word (or letter) was 1500 ms. The duration of each trial was 3000 ms. Subjects indicated their response with bilateral button presses on a four-button keypad, and their response times in milliseconds were recorded. If the answer was "yes," they pressed the two inner buttons, and if the answer was "no," they pressed the two outer buttons. They were told to respond as quickly as possible without sacrificing accuracy, and they were told that they must respond within the 1500 ms time window during which the property word remained on the screen. It is important to note that subjects were *not* given explicit instructions on *how*

to make their decision (e.g., no imagery instructions were given); rather they were told to determine as accurately and as quickly as possible whether the property was part of the concept. Four sample items, using nontest pairs, were shown to the subjects to familiarise them with the task prior to the scanning session.

The timing and trial composition of the blocks were the same for both tasks. Each experimental run consisted of 20 blocks of trials, alternating between property verification blocks and letter verification blocks. True and false trials were intermixed within each property verification block (see Figure 1). There were 10 blocks (i.e., 100 trials) of trials in each run. There were 50 property verification trials and 50 letter verification trials. Among all of the trials within each run, half of them required a “yes” response, and half of them required a “no” response. A total of four runs (i.e., 400 trials) comprised the experiment.

Using PsyScope software (Cohen, MacWhinney, Flatt, & Provost, 1993), stimuli were presented by a Macintosh G3 Powerbook connected to an LCD projector. The image was projected to a screen that was placed at the subjects’ feet. A mirror that was mounted to the head coil was placed right above the subjects’ eyes, and it allowed them to see the projected image on the

screen. However, a different stimulus presentation system was used for three of the subjects (one in the associated group and two in the unassociated group). For these three subjects, the Avotec Silent Vision Visual Presentation system was used (Avotec, 2002). A four-button fibre-optic response pad connected to the computer was used to record the subjects’ button press responses.

Image acquisition

Following the acquisition of sagittal and axial T1-weighted localiser images, gradient echo, echoplanar fMRI was performed in 21 contiguous 5 mm axial slices (TR = 2000, TE = 50, 64 × 64 pixels in a 24 cm field of view, voxel size = 3.75 mm × 3.75 mm × 5 mm) using a 1.5-T GE Signa system equipped with a fast gradient system and the standard quadrature head coil. To minimise head motion, foam padding was placed between the subject’s head and the head coil. Twenty seconds of “dummy” gradient and rf pulse preceded the actual data acquisition to approach steady-state magnetisation.

Image processing

Offline data processing was performed using VoxBo software (VoxBo, 2002). After image reconstruction, the data were sinc interpolated in time to correct for the fMRI acquisition sequence. A slicewise motion compensation method was utilised to remove spatially coherent signal changes by the application of a partial correlation method to each slice in time (Zarahn, Aguirre, & D’Esposito, 1997). Additional motion detection and correction was undertaken using a six-parameter, rigid-body transformation. None of the subjects had translational motion that exceeded 2 mm in any plane or angular motion that resulted in more than a 2 mm displacement. Additionally, spatial smoothing and normalisation were performed. Raw data for all runs from each subject were transformed to standardised MNI space (Evans, Collins, Mills, Brown, Kelly, & Peters, 1993) and spatially smoothed by convolution with a three-dimensional Gaussian kernel that has a FWHM of 1.5 × 1.5 × 2.0 (in voxels).

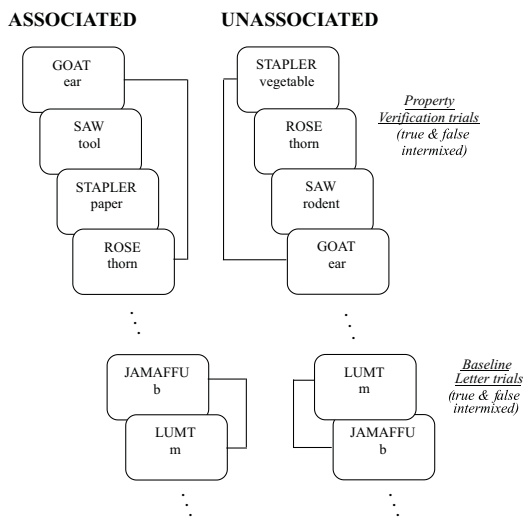


Figure 1. Pictorial depiction of trial types.

Image analyses

Voxelwise analysis was performed on each subject's data by using a general linear model for serially correlated error terms (Worsley & Friston, 1995), and an estimate of intrinsic temporal autocorrelation was included within the model (Aguirre, Zarahn, & D'Esposito, 1997). Furthermore, sine and cosine regressors for frequencies below that of the task (0.0148 Hz) were also included in the general linear model. Temporal data were smoothed with an empirically derived estimate of the haemodynamic response of the fMRI system; this analysis has been empirically demonstrated to hold the mapwise false positive rate at or below tabular values (Zarahn et al., 1997).

RESULTS

Behavioural results

The task employed in this experiment was a direct replication of Solomon's (1997) behavioural experiments. A summary of the relevant data from Solomon's study is reported in Table 1 (for a full report of the behavioural effects of these manipulations, please refer to Solomon's study). In this experiment, we collected and analysed behavioural data from nine of our subjects as a manipulation check. Overall, we replicated the basic findings reported by Solomon: Subjects were highly accurate at verifying properties for concepts and for verifying letters in nonwords. Specifically, individuals in the associated condition ($M = 98.0\%$) were equally accurate at letter verification as subjects in the unas-

sociated condition ($M = 97.8\%$). However, subjects in the associated condition ($M = 93.7\%$) were less accurate at verifying properties than subjects in the unassociated condition ($M = 99.8\%$). In terms of reaction times, on correct true trials, subjects in the associated condition responded slightly slower on the true trials ($M = 829$ ms) than subjects in the unassociated condition ($M = 809$ ms). Furthermore, on correct false trials, subjects in the associated condition responded slower on the associated false trials ($M = 885$ ms) than subjects in the unassociated condition on the unassociated false trials ($M = 833$ ms).

fMRI results

The same analysis was performed on data from both the associated and the unassociated groups: To identify a significant property verification main effect, all property verification trials were compared to all baseline letter verification trials. In both conditions, all subjects showed significant increases in fMRI activity during the property verification condition relative to the baseline condition, mapwise $\alpha = .05$ (see Figure 2). Two analyses were performed to examine group differences in the pattern of activation between subjects in the associated and unassociated groups: First, we performed a focused region of interest analysis to test for group differences in the specific area of visual cortex hypothesised to be involved in imagery and visual semantics. Second, we conducted an exploratory analysis across the whole brain to identify other candidate regions displaying group differences, in either direction.

Table 1. Summary of behavioural results (mean reaction times and mean error rates) from Solomon's (1997) study and the current study

Conditions	Mean reaction times (SD)				Mean error rate			
	Solomon ($n = 96$)		Current study ($n = 9$)		Solomon ($n = 96$)		Current study ($n = 9$)	
	True	False	True	False	True	False	True	False
Associated	842 (182)	979 (276)	829	885	0.07	0.07	0.06	0.07
Unassociated	731 (169)	748 (150)	809	833	0.05	0.02	0.01	0.01

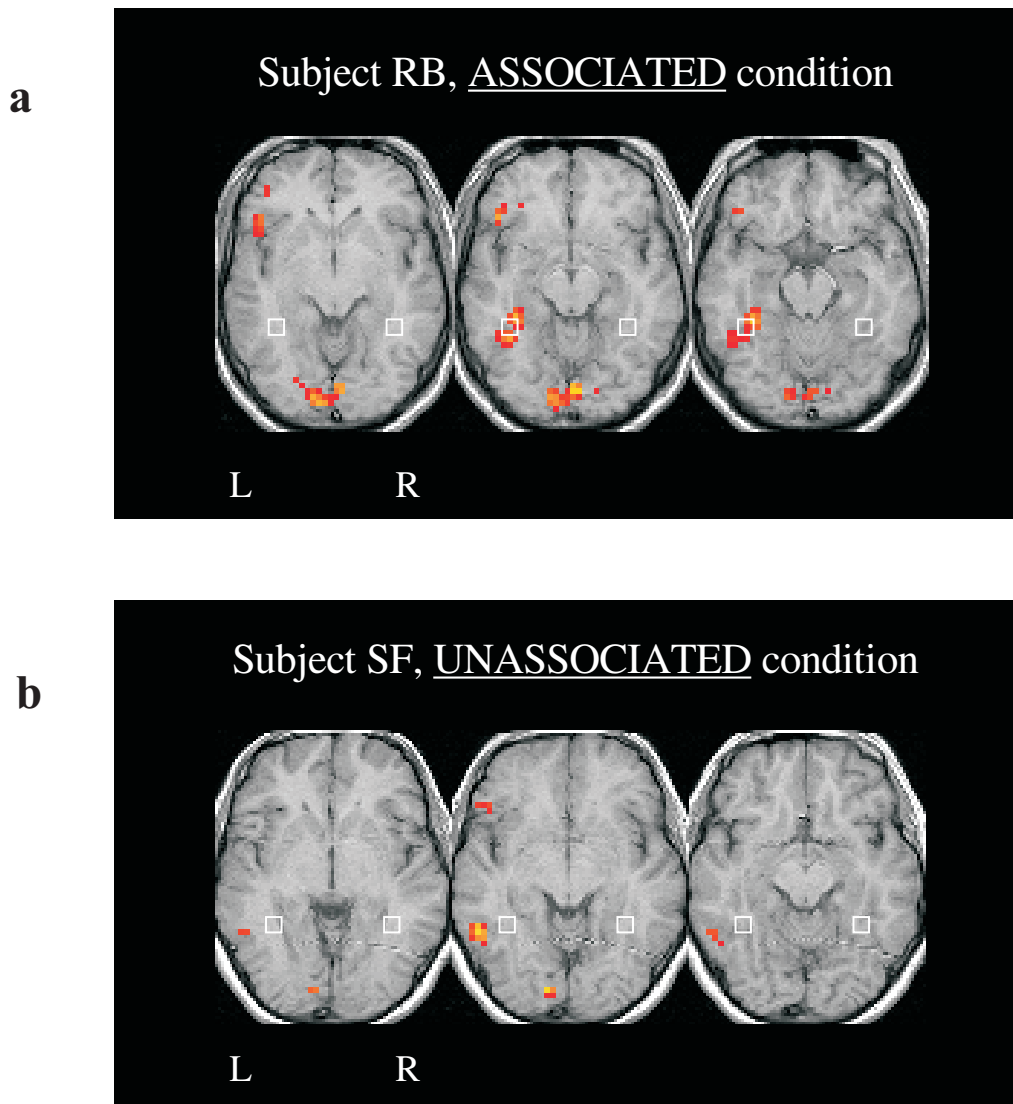


Figure 2. Data from representative subjects in (a) the associated condition and (b) the unassociated condition. Both subjects showed significant increases in fMRI activity during property verification condition relative to baseline condition (mapwise $\alpha = .05$). However, only the associated subject showed significant activation in the left fusiform gyrus. Furthermore, neither subject showed significant activation in the right fusiform gyrus. Anatomically defined regions of interest in the left and right fusiform gyri are indicated by the white outline boxes.

ROI analyses. In each subject, an anatomically defined region of interest (ROI) in the left fusiform gyrus was created. Specifically, the local maxima coordinates for the left fusiform activity reported by D'Esposito et al. (1997) were used as a landmark (Talairach coordinates: $-33, -48, -18$). Subse-

quently, using those coordinates as the centre, a sphere that encompassed 30 voxels ($\sim 2.11 \text{ cm}^3$) was created. This ROI was then used to examine the effects of property verification. The time series from all 30 voxels were averaged, and the effect size for property verification relative to baseline trials




Figure 3. Average t -value of property verification minus baseline letter verification comparison within the left fusiform gyrus and the right fusiform gyrus.

was calculated for the averaged time series for each subject. An unpaired t -test of these effect sizes, with subjects as a random variable and condition as a between-subject variable, was used to test for group differences. Furthermore, to examine potential hemispheric differences, the same approach was used to analyse activity within the right fusiform gyrus (Figure 2).

As a group, subjects in the associated condition showed significant activation in the left fusiform gyrus, $M = 3.09$; $t(6) = 3.05$, $p < .05$ (Figure 3). On the other hand, subjects in the unassociated group did not show significant activity in the same anatomical region $M = 0.26$; $t(6) = 0.451$, $p = .67$

(Figure 3). An unpaired t -test revealed a significant difference between the two groups, $M = 2.83$; $t(12) = 2.44$, $p < .05$.

To investigate laterality differences, a 2 (condition) \times 2 (hemispheres) mixed ANOVA was conducted. A significant interaction was found, $F(1, 12) = 4.80$, $p = .05$. Probing with simple effects revealed a significant hemispheric difference in the associated group, $F(1, 6) = 10.27$, $p < .05$, but not in the unassociated group, $F(1, 6) = 3.85$, $p = .10$.

Whole-brain analyses. We performed an exploratory whole-brain analysis that directly compared all voxels in the associated and the unassociated conditions, using a random effects group analysis of normalised data from all subjects. Due to the small number of subjects in each group, we did not have enough power to detect above-threshold differences using the conservative Bonferroni correction for $\sim 18,000$ voxels. We identified all voxels above a threshold of $t = 3.5$, based on an uncorrected alpha level of .005. To guard against Type I error, we initially included only activation with an extent of seven or more contiguous voxels. Using these criteria, the only area found to be different between the two conditions, with greater activation in the associated condition than the unassociated condition, was the left fusiform gyrus (see Figure 4). This region was slightly superior (~ 10 mm) to the hypothesised region in the ROI analysis. When the cluster size was lowered to two contiguous voxels, the left middle frontal gyrus also revealed greater activation for the associated condition. Furthermore, right middle temporal gyrus showed greater activation for the unassociated condition than the associated condition (see Table 2).

Table 2. Local maxima for areas of activation in the random effects group analysis of the whole brain

Region (Brodmann's area)	Talairach coordinates (x, y, z)			Cluster size ^a	t	
<i>Associated > Unassociated</i>						
Fusiform gyrus (37)	L	-45	-53	-6	9	4.33
Middle frontal gyrus (10)	L	-33	55	-6	2	4.05
<i>Unassociated > Associated</i>						
Middle temporal gyrus (21)	R	48	-24	-6	2	3.81

^aNo. of voxels.




Figure 4. Random effects group analysis of the whole brain. (Associated-Unassociated, cluster size = 7 contiguous voxels.)

DISCUSSION

In this study, we investigated the involvement of visual cortex during conceptual processing, in order to test claims about the representation of semantic knowledge. Specifically, we investigated the extent to which the left occipito-temporal region of the brain is involved in the property verification task. We examined a region of interest in the left fusiform gyrus, which has been associated with mental imagery in previous studies (D'Esposito et al., 1997; Thompson-Schill et al., 1999). We hypothesised that if semantic knowledge is indeed grounded in the perceptual system, we should find significant activity in the left fusiform gyrus during property verification, but only when conceptual information is required (i.e., when associated false trials were present). Two aspects of the results of this study are noteworthy: The location of activation associated with the property verification task and the conditions under which this activation was observed. Each of these points is elaborated below.

The first important finding of this study was that retrieval of semantic knowledge, in the absence of explicit mental imagery instructions, activated a region of visual association cortex that has been associated with visual object recognition (e.g., Ishai, Ungerleider, Martin, & Haxby, 2000; for a review, see Kanwisher, Downing, Epstein, & Kourtzi, 2001) and visual imagery (e.g., D'Esposito et al., 1997). Activation of visual association cortex during semantic processing supports the idea that conceptual knowledge is grounded in the perceptual system. Furthermore, we replicated the hemispheric asymmetry reported in other studies on mental imagery (e.g., D'Esposito et al., 1997; Farah, 1984; Riddoch, 1990). The absence of activation in the right fusiform gyrus serves as a control region to show that effects reported here are specific to areas involved in mental imagery rather than nonspecific increases in activation related to task difficulty or duration of processing. In addition, the hemispheric differences may have implications for understanding the processes involved in retrieval of semantic knowledge. Farah (1995) has claimed that the process of mental image formation is lateralised to the left hemisphere, despite the fact that visual

perception occurs, at least in its initial stages, symmetrically in both hemispheres. Barsalou (1999) argued that retrieval of conceptual knowledge requires the simulation of perceptual information; our results suggest that the process of visual simulation during semantic retrieval is, like mental imagery, a left-hemisphere lateralised process. Differential involvement of the left and right fusiform gyri has also been reported in studies of object recognition. In an fMRI experiment using a repetition priming paradigm, Koutstaal, Wagner, Rotte, Maril, Buckner, and Schacter (2001) reported differential neural effects of priming in left and right fusiform gyri. Specifically, it was reported that visual form changes between first and second exposures were more "costly" in the right fusiform than in the left fusiform. That is, perceptual changes of stimuli led to a greater reduction of neural priming effects in the right fusiform than the left fusiform. Future studies may help elucidate the nature of hemispheric specialisation in the fusiform gyrus, as it relates to both image formation and object recognition.

The second important finding was that activation of this region only occurred under conditions hypothesised to require conceptual knowledge. Within the left fusiform gyrus, we found reliable activity across subjects when true trials were intermixed with associated false trials, but not when true trials were intermixed with unassociated false trials. This pattern of activation follows from the task analysis provided by Solomon (1997). Based on patterns of response times, she argued that two processes are available to perform the property verification task: First, one can consult information about lexical co-occurrence, or strength of association between two words; this process will only be effective when associative strength is a reliable indicator of the correct response (i.e., in the unassociated condition). Second, one can retrieve semantic information about the properties of an object by perceptually simulating the representation; this process is assumed to be less efficient than the former, and will thus be preferred only when the other process is ineffective (i.e., in the associated condition). According to this framework, differences in activation between the associated and

unassociated condition will reflect the semantic retrieval and perceptual simulation demands that are unique to the associated condition. As predicted, these differences were observed in visual association cortex.

It is possible that activation in other regions of cortex would also differ between the two conditions, either as a result of an extended network of cortex that supports perceptual simulation or as a consequence of other cognitive differences between the two tasks (e.g., conflict monitoring in the more difficult, associated condition). In order to evaluate other, unanticipated differences, we conducted an exploratory analysis of effects across the entire brain. With a cluster size of seven contiguous voxels, the left fusiform gyrus was the only region that revealed greater activity in the associated than the unassociated condition. This result is consistent with the region of interest analysis reported above, although the location of the cluster is slightly superior to the *a priori* region of interest based on the results of D'Esposito et al. (1997). Although the results of the whole-brain analysis must be interpreted with caution because a conservative threshold was not applied, the trend in the data indicates that the left fusiform gyrus is selectively recruited by whatever processing differences exist between the two conditions. Based on previous research, differences in the left fusiform gyrus are likely to reflect perceptual simulation (or image formation) that accompanies semantic knowledge retrieval.

The task analysis described above would seem to support the prediction that regions associated with retrieval of lexical associations would be more active in the unassociated condition relative to the associated condition. In the unassociated minus associated contrast, the only area that showed a hint of activation was in the right middle temporal gyrus. However, note that this activation is present only when the cluster threshold is lowered to two contiguous voxels (see Table 2), so it would be premature to draw any strong inferences based on this result. One possible reason for the lack of greater activation in the unassociated condition may be due to automatic word association processing, which is common to both conditions. According to the

priming literature, word association is assumed to be a relatively automatic process (e.g., Meyer & Schvaneveldt, 1971; Shelton & Martin, 1992). When subjects perform the task, word association processes may be triggered automatically. Thus, brain areas that are recruited for word associations would be activated for subjects in both associated and unassociated conditions, and a direct comparison would reveal no differences in relevant brain regions. Another possible explanation for the lack of increased activation in the unassociated condition is that perhaps there is more between-subject variability in the patterns of activation associated with retrieval of linguistic information. Since a group analysis was performed, only areas that are reliably activated for all subjects were identified. Therefore, increased between-subject variability associated with a given process would result in fewer common areas found to be active in the group analysis.

Thus far, we have illustrated how one manipulation (i.e., relationship between true and false trials in a property verification task) modulates activity in the visual association cortex. It is reasonable to further examine how other factors may affect activity in this part of the brain. Data from neuroimaging (e.g., Martin, Wiggs, Ungerleider, & Haxby, 1996; Mummery, Patterson, Hodges, & Wise, 1996; Thompson-Schill et al., 1999) and neuropsychological studies (e.g., Farah, McMullen, & Meyer, 1991; Warrington & McCarthy, 1987; Warrington & Shallice, 1984) have indicated that there may be category-specific effects in the retrieval of semantic knowledge. It is, then, reasonable to suppose that the effects that we have observed in the current study may differ across different types of stimuli. For example, perhaps the processes involved in verifying properties of living things may differ from verifying properties of nonliving things. Unfortunately, since a blocked design was used, we are unable to address that issue in this study. Thus, it will be informative to investigate these effects in an imaging study using an event-related design.

For the past few decades, Pylyshyn has maintained that subjects in mental imagery experiments

are merely simulating the use of visual spatial representations, even though the internal representations are organised amodally (for a review, see Pylyshyn, 1981). He further argued that this is particularly true in experiments that explicitly ask subjects to imagine a certain object or scenario. He postulated that the subjective experience of “seeing with the mind’s eye” is purely epiphenomenal. Researchers who support the idea that conceptual knowledge is grounded perceptually have tried to use neuroimaging data to counter Pylyshyn’s argument. For example, in a review chapter, Kosslyn and Thompson (2000) concluded that visual perception and mental imagery share many common neural mechanisms on the basis of many neuroimaging studies. Although it is difficult to imagine a way in which subjects could “simulate” blood flow to certain regions of the brain, Pylyshyn (in press) has recently extended his criticisms to the functional neuroimaging literature as well.

We believe that our current study addresses part of Pylyshyn’s concern. In the present experiment, although subjects were told to verify each property with the corresponding concept, they were *not* given explicit imagery (or nonimagery) instructions on how to perform the task. Thus, it seems improbable that subjects would interpret the task as requiring them to simulate the use of visual representations. Furthermore, it is even less plausible that all of the associated subjects would choose to simulate the use of visual spatial representations and that the unassociated subjects would choose not to simulate that experience. Our finding that manipulating the type of false trials modulates activation of the visual association areas presents a strong argument against Pylyshyn’s (in press) claims. If neural activity is purely epiphenomenal, there is no reason to expect the group differences observed in this experiment. It is difficult to explain the group differences observed in this study with Pylyshyn’s reasoning. In contrast, our results are easily explained by differences between the two conditions with regard to requirements for conceptual knowledge, which is grounded in the perceptual system.

REFERENCES

- Aguirre, G. K., Zarahn, E., & D’Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. II. Spatially smoothed data collected under null-hypothesis and experimental conditions. *NeuroImage*, *5*, 199–212.
- 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.
- Avotec. (2002). Retrieved February 8, 2002, from <http://www.avotec.org>
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–660.
- Barsalou, L. W., Solomon, K. O., & Wu, L. L. (1999). *Perceptual simulation in conceptual tasks*. Paper presented at the Proceedings of the 4th conference of the International Cognitive Linguistics Association, Amsterdam, The Netherlands.
- Bisiach, E., & Luzzatti, C. (1978). Unilateral neglect of representational space. *Cortex*, *14*, 129–133.
- Charlot, V., Tzourio, M., Zilbovicius, M., Mazoyer, B., & Denis, M. (1992). Different mental imagery abilities result in different regional cerebral blood flow activation patterns during cognitive tasks. *Neuropsychologia*, *30*, 565–580.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers*, *25*, 257–271.
- D’Esposito, M., Detre, J. A., Aguirre, G. K., Stallcup, D., Alsop, D. C., Tippet, L. J., & Farah, M. J. (1997). A functional MRI study of mental image generation. *Neuropsychologia*, *35*, 725–730.
- Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L., & Peters, T. M. (1993). *3D statistical neuroanatomical models from 305 MRI volumes*. Paper presented at the IEEE-Nuclear Science Symposium and Medical Imaging Conference.
- Farah, M. J. (1984). The neurological basis of mental imagery: A componential analysis. *Cognition*, *18*, 245–272.
- Farah, M. J. (1995). Current issues in the neuropsychology of image generation. *Neuropsychologia*, *33*, 1455–1471.
- Farah, M. J., McMullen, P. A., & Meyer, M. M. (1991). Can recognition of living things be selectively impaired? *Neuropsychologia*, *29*, 185–193.

- Farah, M. J., Soso, M. J., & Dasheiff, R. M. (1992). Visual angle of the mind's eye before and after unilateral occipital lobectomy. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 241–246.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.
- Gentner, D. (1978). On relational meaning: The acquisition of verb meaning. *Child Development*, *49*, 988–998.
- Glaser, W. R. (1992). Picture naming. *Cognition*, *42*, 61–106.
- Goldenberg, G., Podreka, I., Steiner, M., Franzen, P., & Deecke, L. (1991). Contributions of occipital and temporal brain regions to visual and acoustic imagery—a SPECT study. *Neuropsychologia*, *29*, 695–702.
- Goldenberg, G., Podreka, I., Steiner, M., & Willmes, K. (1987). Patterns of regional cerebral blood flow related to memorizing of high and low imagery words: An emission computer tomography study. *Neuropsychologia*, *25*, 473–485.
- Intons-Peterson, M. J. (1983). Imagery paradigms: How vulnerable are they to experimenters' expectations? *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 394–412.
- Ishai, A., Ungerleider, L. G., Martin, A., & Haxby, J. V. (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience*, *12*(suppl. 2), 35–51.
- Kanwisher, N., Downing, P., Epstein, R., & Kourtzi, Z. (2001). Functional neuroimaging of human visual recognition. In R. Cabeza & A. Kingstone (Eds.), *The handbook of functional neuroimaging* (pp. 109–152). Cambridge, MA: MIT Press.
- Kosslyn, S. M. (1975). Information representation in visual images. *Cognitive Psychology*, *7*, 341–370.
- Kosslyn, S. M. (1976). Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval time. *Memory and Cognition*, *4*, 291–297.
- Kosslyn, S. M. (1994). *Image and brain*. Cambridge, MA: MIT Press.
- Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 47–60.
- Kosslyn, S. M., & Thompson, W. L. (2000). Shared mechanisms in visual imagery and visual perception: Insights from cognitive neuroscience. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences*. Cambridge, MA: MIT Press.
- Koutstaal, W., Wagner, A. D., Rotte, M., Maril, A., Buckner, R. L., & Schacter, D. L. (2001). Perceptual specificity in visual object priming: Functional magnetic resonance imaging evidence for a laterality difference in fusiform cortex. *Neuropsychologia*, *39*, 184–199.
- Locke, J. (1959). *An essay concerning human understanding, Vol. 1 & 2* (1st ed.) (original work published in 1960). New York: Dover.
- Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature*, *379*, 649–651.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.
- 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 [published erratum appears in Proc R Soc Lond B Biol Sci 1996 Dec 22;263(1377):1755–6]. *Proceedings of the Royal Society of London - Series B: Biological Sciences*, *263*, 989–995.
- Nelson, K. (1974). Concept, word, and sentence: Interrelations in acquisition and development. *Psychological Review*, *81*, 267–285.
- Price, H. H. (1953). *Thinking and experience*. London: Hutchinson's Universal Library.
- Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, *88*, 16–45.
- Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. W. (in press). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*.
- Riddoch, M. J. (1990). Loss of visual imagery: A generation deficit. *Cognitive Neuropsychology*, *7*, 249–273.
- Roland, P. E., & Friberg, L. (1985). Localization of cortical areas activated by thinking. *Journal of Neurophysiology*, *53*, 1219–1243.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Servos, P., & Goodale, M. A. (1995). Preserved visual imagery in visual form agnosia. *Neuropsychologia*, *33*, 1383–1394.

- Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1191–1210.
- Shepard, R. N., & Cooper, L. A. (1982). *Mental images and their transformations*. New York: Cambridge University Press.
- Solomon, K. O. (1997). *The spontaneous use of perceptual representations during conceptual processing*. Chicago: University of Chicago.
- Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory: Inferring “how” from “where.” *Neuropsychologia*, *43*, 280–292.
- Thompson-Schill, S. L., Aguirre, G. K., D’Esposito, M., & Farah, M. J. (1999). A neural basis for category and modality specificity in semantic knowledge. *Neuropsychologia*, *37*, 671–676.
- VoxBo. (2002). Retrieved February 8, 2002, from <http://www.voxbo.org>
- Warrington, E. K., & McCarthy, R. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, *110*, 1273–1296.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829–854.
- Worsley, K. J., & Friston, K. (1995). Analysis of fMRI time-series revisited—again. *NeuroImage*, *2*, 173–182.
- Zarahn, E., Aguirre, G. K., & D’Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *NeuroImage*, *5*, 179–197.
- Zatorre, R. J., Halpern, A. R., Perry, D. W., Meyer, E., & Evans, A. C. (1996). Hearing in the mind’s ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience*, *8*, 29–46.

