What Brain Imaging Can Tell Us About Embodied Meaning

Marcel Adam Just
Center for Cognitive Brain Imaging
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA

Abstract: Brain imaging studies of language processing (using fMRI) can indicate under what circumstances the embodied aspects of language representations become activated. In particular, the processing of language is distributed across a number of cortical centers, including not only classic language areas in association cortex (which might be involved in symbolic processing), but also sensory and motor areas. A set of fMRI studies on visual imagery in sentence comprehension reveals both the perceptual-motor and the symbolic aspects of brain function that underlie language processing. Moreover, they indicate some of the conditions under which perceptual or motor representations are most likely to be activated. Another set of studies on word comprehension indicates that the neural signature of certain concrete semantic categories (tools and dwellings) and individual category exemplars can be identified by machine learning algorithms operating on the fMRI data, and that perceptual and motor representations constitute part of the signature.

Visual imagery in sentence comprehension

Many types of thinking, in particular language comprehension, entail the use of mental imagery. Understanding a text on architecture or automobile design seems impossible without mental imagery. Language often refers to perceptually-based information. For example, to evaluate a sentence like *The number eight when rotated 90 degrees looks like a pair of spectacles*, a reader must process the content of the sentence, retrieve a mental image of the shape of the digit 8, mentally apply a rotation transformation to it, and then evaluate the resulting image. In this case, there seems little doubt that a perceptually-based representation is involved in the comprehension. This perceptually-based representation has sometimes been called the referential representation or the situation model. Our fMRI studies attempt to determine the characteristics of such representations and the conditions under which they are likely to be evoked or activated.

Previous studies have indicated that mental imagery generated by verbal instructions and by visual encoding activate similar cortical regions (Mellet et al., 1996, 1998, & 2002; Mazoyer et al., 2002). Several studies examining mental imagery have observed activation of the parietal area (Just et al., 2001; Mellet et al., 1996 & 2000; Deiber et al., 1998; Ishai et al., 2000, Kosslyn et al., 1993, 1996 & 1999), particularly around the intra-parietal sulcus. Our imagery studies attempted to determine the conditions under which such activation occurs during language comprehension.

There is also a possibility that the neural activity underlying the imagery in language processing is affected by the presentation modality of the language (i.e., written vs. spoken). For example, the neural activity elicited in primary visual cortex during mental imagery following verbal vs. visual encoding was different (Mellet et al., 2000); there was less primary visual activation during imagery after visual encoding compared to verbal encoding, suggesting that presentation modality may indeed affect later imagery processing. Another study by Eddy and Glass (1981) examined how the visual processes in reading might be related to the visual imagery processes that a sentence engenders, comparing visual and auditory sentence presentation modes. High-imagery sentences took longer to verify as true or false than low-imagery sentences when the sentences were presented visually, but not when they were presented auditorally. These findings again suggest that the presentation modality of

a sentence may affect the processing of the subsequent imagery.

Our imagery studies examined mental imagery processes in the context of a language comprehension task (Just et al., 2004). One of the main goals was to examine the interaction between two somewhat separable neural systems, the mental imagery and language processing systems. In the context of the embodiment debate, these studies ask not whether embodied (perceptual or motor) activation occurs, but the circumstances under which it occurs and how it is related to other more symbolic activation. To accomplish this goal, we used fMRI to measure not only the activation levels but also the functional connectivities of the regions believed to be involved in mental imagery, to determine the relations between the two systems. A second goal was to examine the effect of input modality, comparing the effect on the imagery-related activation when the sentences were either heard or read.

The study examined brain activation while participants read or listened to high-imagery sentences like *The number eight when rotated 90 degrees looks like a pair of spectacles* or low-imagery sentences, and judged them as true or false. They included sentences requiring various types of spatial transformation or spatial assessment such as mental rotation (like the *spectacles* sentence), evaluation of spatial relations (e.g., *On a map, Nevada is to the right of California*), combination of shapes (e.g., *The number nine can be constructed from a circle and a horizontal line*, a false example), and comparison of visual aspects of common objects (e.g., *In terms of diameter, a quarter is larger than a nickel, which is larger than a dime*). Although these sentences generally required that a spatial transformation be mentally performed, pilot studies indicated that understanding a complex spatial description without a transformation produced similar results. The low-imagery sentences could be verified by referring to general knowledge, without the use of imagery (e.g., *Although they are now a sport, marathons started with Greek messengers bringing news*).

The sentence imagery manipulation affected the activation in regions (particularly the left intraparietal sulcus) that activate in other mental imagery tasks, such as mental rotation. Both the auditory and visual and auditory presentation experiments indicated much more activation of the intraparietal sulcus area in the high-imagery condition as shown in Figure 1, suggesting a common neural substrate for language-
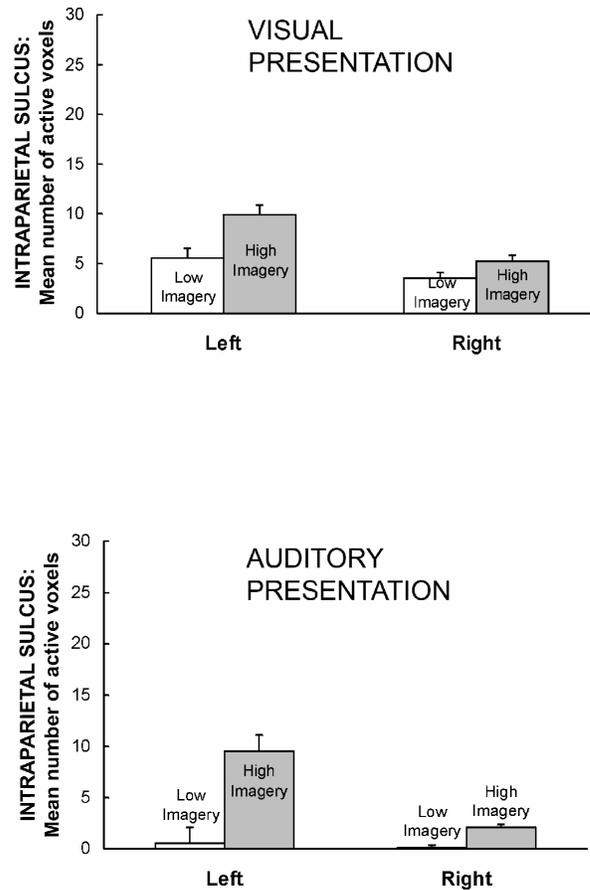

Figure 1. The activation in the intraparietal sulcus area is higher for high imagery sentences (particularly in the left hemisphere), and the effect is similar regardless of whether the sentences are presented visually or auditorally.

evoked imagery that is independent of the input modality. There was more activation in the intraparietal sulcus area in the reading than in the listening condition (probably owing to the attentional demands of *directing spatial attention and possibly eye movements to particular sentence locations during* reading), but the magnitude of the imagery effect was comparable in the two cases.

Functional connectivity and imagery. The various anatomical regions of the cortex involved in processing a task must be able to effectively communicate and synchronize their processes for the system to function. In a language task, this means that the areas responsible for executing subcomponent processes must collaborate to synthesize the information

In M. de Vega, A. Glenberg, & A. Graesser (Eds.), 2008   Neural embodiment of sentence and word meaning
Symbols and Embodiment: Debates on meaning and cognition (pp. 75-84)
Oxford, UK: Oxford University Press.

necessary for comprehension. Such collaboration can be measured in functional neuroimaging studies by computing the correlation of the activation time series in a given region with the activation time series of another region. The extent to which the activation levels of two regions rise and fall in tandem is taken as a reflection of the degree to which the two regions are functionally connected, and the term that is widely used to refer to the activation time series correlation is *functional connectivity*. Previous research has provided some evidence that as task demands increase, functional connectivity also increases, for example, as a function of working memory load (e.g., Diwadkar, Carpenter, & Just, 2000), reflecting the need for tighter coordination in a more demanding condition. Functional connectivity has also been shown to be modulated by comprehension difficulty, and differentially so for people of different working memory capacity (Prat et al., 2007). A relation between functional and anatomical connectivity has been demonstrated in autism, where the functional connectivity between cortical regions is correlated with the size of the corpus callosum segment that anatomically connects them (Just et al., 2007).

In addition to exhibiting higher activation levels during the processing of high-imagery sentences, the left intraparietal sulcus also showed greater functional connectivity in this condition with other cortical regions, particularly with language processing regions, regardless of the input modality. The imagery manipulation affected the functional connectivity between the left intraparietal sulcus area and other brain areas that are activated in this task. It also the affected the functional connectivity between the left superior temporal (Wernicke's) area and other brain areas. These two areas are proposed to be respectively involved in the imagery processing of high-imagery sentences (left intraparietal) and semantic (symbolic) processing (Wernicke's). The five other activated brain areas were left hemisphere areas, including areas centrally involved in language processing (pars opercularis, pars triangularis, dorsolateral prefrontal cortex, frontal eye fields, and inferior parietal lobule). The intraparietal sulcus area had higher functional connectivities with the five activated brain areas when the sentences were high in imagery, whereas the left temporal area had higher functional connectivities to these five areas for low-imagery sentences. This result applies to be the visual and auditory presentation conditions, as shown in Figure 2.
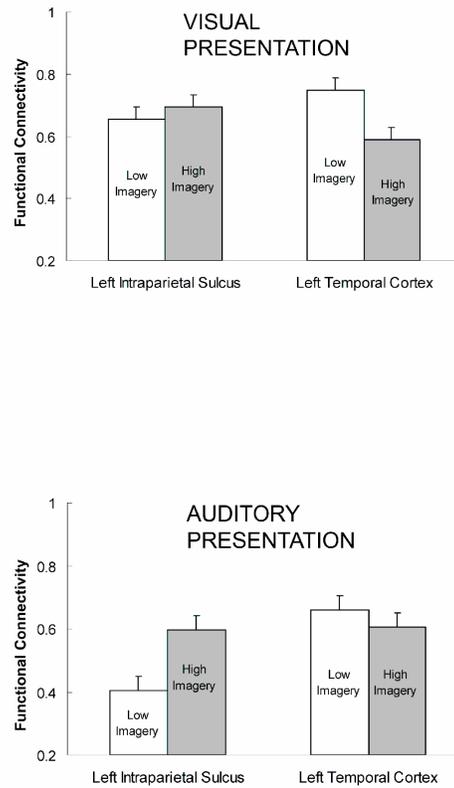

Figure 2. The average functional connectivity of the left intraparietal sulcus ROI (LIPS) and the left temporal ROI (LT) and with five left hemisphere ROIs: pars opercularis, pars triangularis, the inferior parietal lobule, dorsolateral prefrontal cortex, and frontal eye fields. The functional connectivity between LIPS and these areas is greater for high-imagery sentences. The opposite is true for LT, which is not involved in imagery processing. The same pattern occurs for both visual and auditory presentation of sentences.

The result provides important converging evidence for implicating the intraparietal sulcus in imagery processing in sentence comprehension, and more directly indicates the higher degree of functional interaction between this embodied imagery representation and some of the other key activated regions in the high imagery condition.

Sentence imagery in autism. In a recent study, we examined the processing of such high- and low-

In M. de Vega, A. Glenberg, & A. Graesser (Eds.), 2008
Symbols and Embodiment: Debates on meaning and cognition (pp. 75-84)
Oxford, UK: Oxford University Press.

Neural embodiment of sentence and word meaning

imagery sentences in adults with high-functioning autism, in comparison to age- and IQ-matched controls. This study provides the opportunity to determine whether the use of imagery (or embodied representations) in sentence comparison might be disrupted in a special neurological population. Can meaning embodiment be disrupted or modulated by a neurological disorder?

First, consider the effect of imagery on the control participants, which is very similar to the study described above. Figure 3 shows the imagery effect (high imagery minus low), displaying the prominent extra activation in the parietal area, as well as some in the prefrontal and inferior temporal areas.
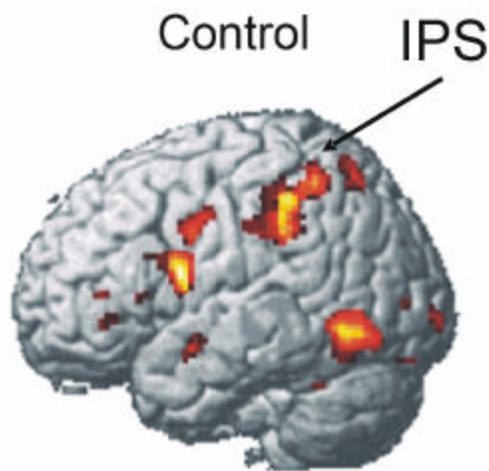
high-functioning autism, entitled her book *Thinking in Pictures*. In an fMRI letter n-back study, Koshino et al. (2005) found of more visual coding of letters in autism compared to a verbal coding strategy in the controls. Similar results were also found in a faces working memory task (Koshino et al., in press). These studies indicate that there is a tendency in people with autism to use more visuospatial regions by recruiting posterior brain regions in accomplishing varied tasks, including language tasks.

The group with autism showed similar activation to the controls in the processing of high-imagery sentences (prominent activation in the parietal area, particularly around the intraparietal sulcus). It is
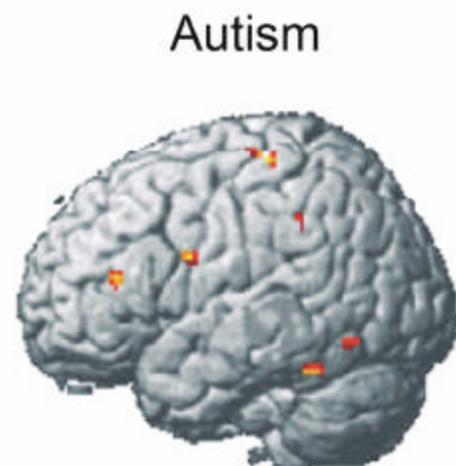

Figure 3. High minus low imagery effect of visual sentence comprehension in control participants, displaying a large imagery effect in parietal, prefrontal, and inferior temporal areas.

Figure 4. High minus low imagery effect of visual sentence comprehension in people with autism, displaying a negligible imagery effect.

What of people with high-functioning autism? There have been frequent suggestions that spatial/perceptual processing is spared (or sometimes even enhanced) in autism. For example, Shah and Frith (1983; 1993) found that participants with autism have an advantage at certain types of spatial tasks. When a task is amenable to either a visual or a verbal strategy, there is a suggestion that people with autism prefer a visual strategy. There are many informal reports that individuals with autism are predominantly visual thinkers (Grandin, 1995). Temple Grandin, an accomplished professional with

reasonable to conclude that in the high-imagery condition, which can be said to make use of embodied (perceptual) representations, the processing of the autism group resembled the processing of the control group. However, the interesting new result was that unlike the control group, the autism group displayed a similar amount of parietal imagery-related activation even in the low-imagery condition. Figure 4 shows the minimal imagery effect (the subtraction of high minus low imagery). It appears that the autism group uses embodied representations even in contexts where control participants do not.

In M. de Vega, A. Glenberg, & A. Graesser (Eds.), 2008   Neural embodiment of sentence and word meaning
Symbols and Embodiment: Debates on meaning and cognition (pp. 75-84)
Oxford, UK: Oxford University Press.

The tentative account for why this occurs is that the cortical connectivity of the neural system is compromised in autism (the main tenet of the underconnectivity theory of autism), particularly affecting the communication between frontal and more posterior areas, resulting in a greater reliance on the more posterior areas. Thus in autism, there is often less frontal activation and more posterior activation in a number of different types of tasks. One might say that cognition is more embodied or concrete in autism, and less abstract.

Visual imagery in the comprehension of novel metaphors. Mason, Eviatar, and Just (under review) compared the comprehension of frozen and novel metaphors, and found an interesting difference that bears on embodiment. During the comprehension of a frozen metaphor passage, the same language processing areas are active as in normal reading (e.g., left middle and superior temporal lobe, left inferior frontal gyrus, dorsolateral prefrontal cortex bilaterally, right middle and superior temporal, and superior medial frontal areas). However, it is the contrast with novel metaphors that is central here. The novel metaphors evoked parietal activation around the intraparietal sulcus area bilaterally, suggesting that visual imagery processes were being used to instantiate and/or interpret the novel metaphors. The novel metaphors were generally visual in nature, such as a metaphor comparing a winding road to a ribbon. These results demonstrate the selective use of imagery in metaphor comprehension, suggesting that perceptual representations are used in the comprehension of novel but not frozen metaphors. It may be that the mappings between domains that are needed in novel metaphor comprehension (e.g. the mapping between the solar system and the atom, to take a clearly spatial example) are often mediated through spatial representations. A more general principle of neural function is that a given representation or process can be activated on an "as-needed" basis. In this view, the embodiment occurs for only those types of metaphors that require it for their appropriate comprehension.

Selectivity of the perceptual representations activated in a haptic imagery task. The imagery-related activation around the intraparietal sulcus can also be evoked in a haptic imagery task. Participants are asked to compare the either the geometric (visual) properties of two objects (e.g., *Which is larger, a pumpkin or a cucumber?*) or haptic (touch) properties (*Which is softer, a lemon or an egg?*) (Newman et al.,

2005). When the haptic properties (e.g. hardness) of the objects are interrogated via haptic imagery, then there is extra activation in the inferior extrastriate region. Additionally, a region in the lateral occipital cortex is activated for either type of probe. Thus, thinking about the object in almost any way evokes *some* perceptual activation in a number of areas, but the amount of activation increases in areas whose information is probed. So there is selectivity in the perceptually-related activation.

Summary. This set of fMRI studies of imagery provides several lessons about embodiment, particularly highlighting the circumstances under which embodied representations are likely to be in play. First, perceptual (embodied) representations are used when there is some mental action or assessment to be performed on the perceptual representation in order to comprehend the sentence or perform the task. Second, in the processing of sentence imagery, the parietal activation associated with the perceptual representation is synchronized with the symbolic language areas with which it is collaborating. Third, in autism, perceptual representations are used for comprehending even low-imagery sentences, which control participants process without the benefit of perceptual representations. And fourth, perceptual representations are used in the comprehension of novel metaphors but not frozen metaphors. All of these lessons have an overarching theme. The embodied representations are activated more often or to a higher level in those situations in which the perceptual information is particularly useful or salient.

The embodied neural signature of words referring to object concepts

In a multidisciplinary project done in collaboration with Tom Mitchell and our two research groups, particularly involving Robert Mason, Svetlana Shinkareva, Vincente Malave, and Francisco Peirera, we have been using machine learning techniques that can learn what activation pattern defines a cognitive state. Our investigation of the representation of words and sentences attempted to determine the neural signature of a semantic category using an engineering rather than a neuroscience approach. We started with concrete nouns as stimuli, taken from a small number of categories. In one of our studies, we showed participants a total of 14 different words, including seven tools (e.g., hammer, hatchet, pliers, screwdriver) and seven dwellings (e.g., mansion, castle, palace, hut, apartment). The experiments and the data analyses are

unconventional, as are the goals. The experimental tasks attempt to isolate individual thoughts of brief duration, typically about 4 sec. We present each word several times during a scanning session (six times in this experiment). We then train the classifiers on all but one of the presentations, and then try to classify the activation patterns in the left-out presentation, iterating through all the possible ways of leaving one presentation out. A classifier, of which logistic regression is a common example, is a decision system which takes as its input the values of some features, and produces as output a discrete label related to the input values. In this application, the input to the classifier is the activation level of a set of voxels (say 200) measured during the reading of a word, and the output is the category of the word that the word belongs to.

In our quest for the neural signature of concepts, we have discovered that the neural representations of at least these concepts contain perceptual and motor information that is pertinent to a category. The embodiment of a category, namely its representation of its perceptual and motor attributes, is part of the category's neural signature. For example, the representations of tools include voxels in the motor area and in the somatosensory area.

There are several ways in which our findings resemble those of two other fMRI projects reported in this volume, in the papers of Pulvermuller. First, the neural signatures span a number of areas distributed over the cortex, including all four lobes in both hemispheres as well as the cerebellum. Second, what differentiates the activation for different categories (what makes it a signature) is not which voxels activate, but the intensity level of their activation.

Our results extend beyond just these commonalities with the other studies. In addition, we have been able to classify with some accuracy the category of the word that is being read. Moreover, we have been able to identify the individual word that a person is reading from the neural signature.

The machine learning techniques were able to identify which of two categories of word is being ready by a participant, with a median rank accuracy over 12 participants of 79% (as of December 2005). (Rank accuracy is the percentile rank of the correct word within a list of predictions made by the classifier. Chance rank accuracy would be 50%). For the participant with the best result, the classifier was extremely accurate in identifying the category of the word being read, with a rank accuracy of 93%.

The classifier was also able to identify to some extent which of the 14 different words was being read, with a median rank accuracy over 12 participants of 69% (still far above chance), and the rank accuracy for the best participant was 76%. The voxels used by the classifier to distinguish among categories and among words are distributed across many regions of the cortex, including the right hemisphere and motor areas. For example, Figure 5 shows a motor region in which several participants had diagnostic voxels, indicating that the motor representation of how a tool is used is part of the meaning of a word that names a hand tool.
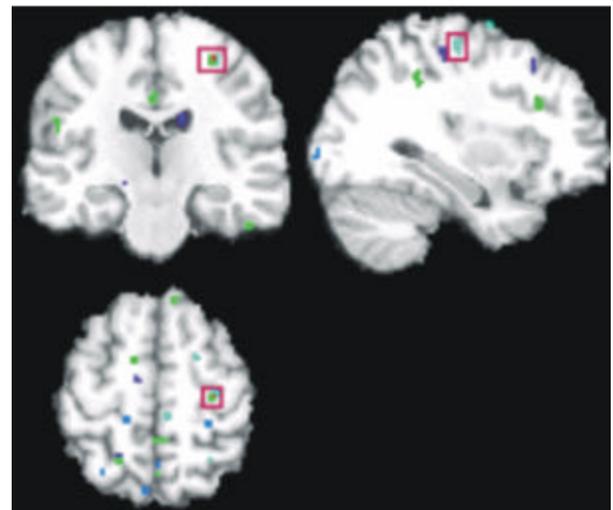

Figure 5. Voxels in the motor region used by the classifier for word identification in several participants.

The rank accuracy of the identification (particularly for distinguishing among the 14 words) is measured as follows. For each test input, the trained classifier outputs a rank-ordered list of the 14 words, ordered from most to least probable according to its learned model. A perfect identification produces a rank accuracy of 1 (the correct answer is ranked as the most likely) and random performance by the classifier produces a rank accuracy of 0.5. Classifier accuracy was measured by the percentile rank of the correct classification in the output sorted list.

The classifier that has been most expedient for our research is a Gaussian Naïve Bayes (GNB) classifier. Other classifiers that we sometimes use include artificial neural networks and logistic regression. The results are somewhat similar across classifiers. The

In M. de Vega, A. Glenberg, & A. Graesser (Eds.), 2008
Symbols and Embodiment: Debates on meaning and cognition (pp. 75-84)
Oxford, UK: Oxford University Press.

Neural embodiment of sentence and word meaning

input to the classifier is represented as a feature vector, where each feature typically corresponds to the observed MR intensity value at a specific voxel at a specific time, relative to a baseline. Although the classification can be done by providing input to the classifier about all of the voxels, the accuracy is generally higher if only a small subset of all the voxels (say 25-200 out of about 15000) is used in the classification. We have used a number of different criteria to select the subset of diagnostic voxels, and most useful criterion has been the consistency of the voxels across presentations in how they respond to the different words.

The neural signature is based on many voxels, some of which are in motor areas, many are in perceptual areas, and some are in frontal (conceptual) areas. Some informal sensitivity analyses were done to determine if some subset of voxels is particularly important or unimportant for accurate classification. So far, no striking systematic differences have been found in how useful the voxels in various locations are. For example, voxels in the motor area are not very different from voxels elsewhere in terms of their contribution to classifying the signature.

Summary. The neural evidence clearly indicates that in at least some cases, perceptual and motor representations are activated during processing that is primarily conceptual. The scientific questions for the future might center around the *when* and the *how* of embodied cognition. Brain imaging has provided a venue for addressing these questions, as well as providing substantiating evidence that embodied cognition is much more prevalent than was generally assumed a decade ago.

Acknowledgments

References

Deiber, M-P, Ibanez, V, Honda, MSN, Raman, R and Hallett, M (1998). Cerebral processes related to visuomotor imagery and generation of simple finger movements studied with positron emission tomography. *NeuroImage, 7*, 73-85.

Diwadkar, VA, Carpenter, PA and Just, MA (2000). Collaborative activity between parietal and dorso-lateral prefrontal cortex in dynamic spatial working memory revealed by fMRI. *NeuroImage*, **12**, 85-99.

Eddy, JK and Glass, AL (1981). Reading and listening to high and low imagery sentences. *Journal of Verbal Learning and Verbal Behavior, 20*, 333-45.

Grandin, T (1995). *Thinking in pictures: And other reports from my life with autism.* New York: Doubleday.

Ishai, A, Ungerleider, LG, Martin, A and Haxby, JV (2000). The representation of objects in the human occipital and temporal cortex. *Journal of Cognitive Neuroscience, 12*, 35-51.

Just, MA, Carpenter, PA, Maguire, M, Diwadkar, VA and McMains, S (2001). Mental rotation of objects retrieved from memory: An fMRI study of spatial processing. *Journal of Experimental Psychology: General, 130*, 493-504.

Just, MA, Cherkassky, VL, Keller, TA, Kana, RK and Minshew, NJ (2007). Functional and anatomical cortical underconnectivity in autism: Evidence from an fMRI study of an executive function task and corpus callosum morphometry. *Cerebral Cortex*, **17**, 951-961.

Just, MA, Newman, SD, Keller, TA, McEleney, A and Carpenter, PA (2004). Imagery in sentence comprehension: An fMRI study. *NeuroImage, 21*, 112-24.

Koshino, H, Carpenter, PA, Minshew, NJ, Cherkassky, VL, Keller, TA and Just, MA (2005). Functional connectivity in an fMRI working memory task in high-functioning autism. *NeuroImage, 24*, 810-21.

Koshino, H, Kana, RK, Keller, TA, Cherkassky, VL, Minshew, NJ and Just, MA. (2008). fMRI investigation of working memory for faces in autism: Visual coding and underconnectivity with frontal areas. *Cerebral Cortex*, 18, 289-300.

Kosslyn, SM, Alpert, NM, Thompson, WL, Maljkovic, V, Weise, SB, Chabris, CF, Hamilton, SE, Rauch, SL and Buonanno, FS (1993). Visual mental imagery activates topographically organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience, 5*, 263-87.

Kosslyn, SM, Pascual-Leone, A, Felician, O, Camposano, S, Keenan, JP, Thompson, WL, Ganis, G, Sukel, KE and Alpert, NM (1999). The role of area 17 in visual imagery: Convergent evidence

In M. de Vega, A. Glenberg, & A. Graesser (Eds.), 2008
Symbols and Embodiment: Debates on meaning and cognition (pp. 75-84)
Oxford, UK: Oxford University Press.

Neural embodiment of sentence and word meaning

from PET and rTMS. *Science,* **284**, 167-70.

Kosslyn, SM., Shin, LM, Thompson, WL, McNally, RJ, Rauch, SL, Pitman, RK and Alpert, NM (1996). Neural effects of visualizing and perceiving aversive stimuli: A PET investigation. *NeuroReport,* **7**, 1569-76.

Mazoyer, B, Tzourio-Mazoyer, N, Mazard, A Denis, M and Mellet, E (2002). Neural bases of image and language interactions. *International Journal of Psychology,* **37**, 204-08.

Mellet, E, Bricogne, S, Crivello, F, Mazoyer, B, Denis, M and Tzourio-Mazoyer, N (2002). Neural basis of mental scanning of a topographic representation built from a text. *Cerebral Cortex,* **12**, 1322-30.

Mellet, E, Bricogne, S, Tzourio-Mazoyer, N, Ghaem, O, Petit, L, Zago, L, Etard, O, Berthoz, A, Mazoyer, B and Denis, M (2000). Neural correlates of topographic mental exploration: The impact of route versus survey perspective learning. *NeuroImage,* **12**, 588-600.

Mellet, E, Petit, L, Mazoyer, B, Denis, M and Tzourio, N (1998). Reopening the mental imagery debate: Lessons from functional anatomy. *NeuroImage,* **8**, 129-39.

Mellet, E, Tzourio, N, Crivello, F, Joliot, M, Denis, M and Mazoyer, B (1996). Functional neuroanatomy of spatial mental imagery generated from verbal instructions. *Journal of Neuroscience,* **16**, 6504-12.

Newman, SD, Klatzky, RL, Legerman, SJ and Just, MA (2005). Imagining material versus geometric properties of objects: An fMRI study. *Cognitive Brain Research,* **23**, 235-46.

Prat, C, Keller, TA and Just, MA (2007). Individual differences in sentence comprehension: An fMRI investigation of syntactic and lexical processing demands. *Journal of Cognitive Neuroscience*, **19,** 1950-1963.

Shah, A and Frith, U (1983). An islet of ability in autistic children: A research note. *J Child Psychol Psychiatry,* **24**, 613-20.

Shah, A and Frith, U (1993). Why do autistic individuals show superior performance on the block design task? *Journal of Child Psychology and Psychiatry,* **34**, 1351-64.