On central cognition

Peter Carruthers

Philosophical Studies
An International Journal for Philosophy
in the Analytic Tradition

ISSN 0031-8116
Volume 170
Number 1

Philos Stud (2014) 170:143-162
DOI 10.1007/s11098-013-0171-1

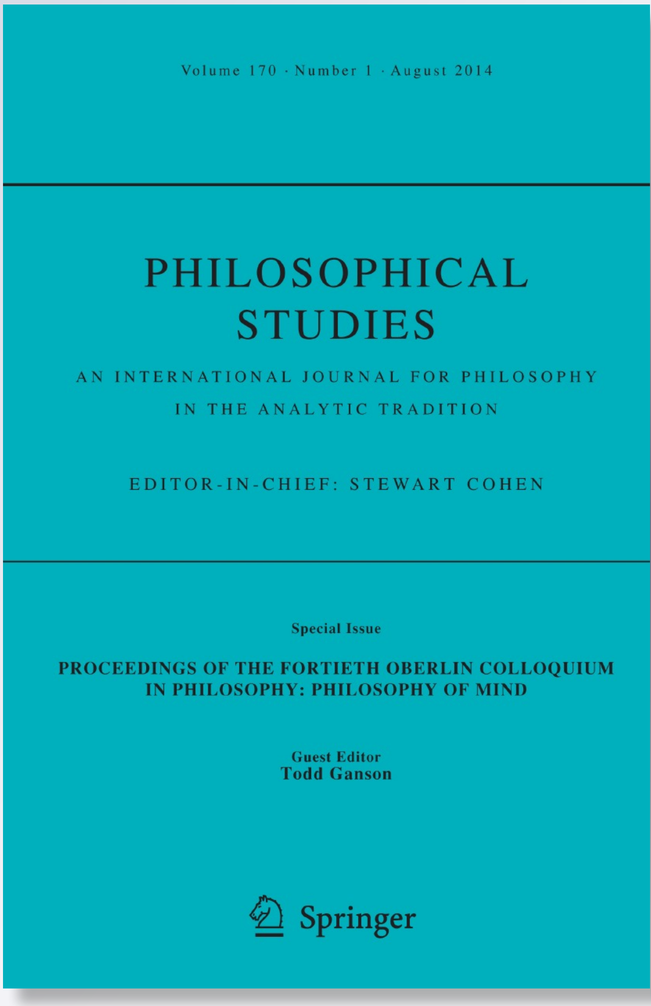

Philos Stud (2014) 170:143–162
DOI 10.1007/s11098-013-0171-1

On central cognition

Peter Carruthers

Published online: 3 August 2013

This article examines what is known about the cognitive science of working memory, and brings the findings to bear in evaluating philosophical accounts of central cognitive processes of thinking and reasoning. It is argued that central cognition is sensory based, depending on the activation and deployment of sensory images of various sorts. Contrary to a broad spectrum of philosophical opinion, the central mind does not contain any workspace within which goals, decisions, intentions, or non-sensory judgments can be active.

1 Introduction: philosophers' commitments

Most philosophers believe that our so-called "personal-level" propositional attitudes—our personal judgments, beliefs, goals, values, decisions, and intentions—are consciously accessible. Many will also allow, of course, that some of the processing that takes place within our perceptual faculties is inaccessible to us; and many will grant the possibility of unconscious attitudes, as well as processes of inference and decision making that are likewise inaccessible to their subjects. But our conscious attitudes are, as such, accessible to us, and they can interact directly with each other and with capacities for inference and decision making to issue in new attitudes of these types, and ultimately in novel forms of behavior. Personal-level attitudes are, it is said, "inferentially promiscuous", in that they can figure in inferences combined with any other such attitudes (Evans 1982; Brewer 1999; Hurley 2006). It is this promiscuity that is said to underlie the distinctive flexibility of human thought and behavior.

P. Carruthers (✉)
Department of Philosophy, University of Maryland, College Park, MD, USA
e-mail: pcarruth@umd.edu

Fodor (1983, 2000) endorses a similar idea. He thinks that any one of our attitudes can in principle be brought to bear in the evaluation of any other, either for purposes of belief-formation or for decision making. Central cognition is, in this sense, said to be *isotropic*. Indeed, Fodor is so confident that the central mind has such properties that he even endorses the notorious pessimistic conclusion which he believes follows from it: for the foreseeable future, we should give up on seeking a science of central cognitive processes of thinking and reasoning.[1]

It appears, then, that many philosophers share a conception of the mind as containing, at its core, a workspace within which any of our conscious propositional attitudes can become active. Many go further, claiming that concepts may be freely combined with one another in this workspace to issue in novel thoughts, hypotheses, or suppositions. Evans (1982), for example, proposes a famous "generality constraint" on genuine concept-possession, which has been widely endorsed by philosophers since (Peacocke 1992; Camp 2004). The constraint is that each of the concepts possessed by a subject should be capable of freely combining with any others of appropriate adicity to form novel thoughts. Hence if one possesses the concepts *F* and *a* needed to entertain the thought *Fa*, and one can also combine the concepts *G* and *b* in the thought *Gb*, then one must be capable of entertaining the thoughts *Fb* and *Ga*. Likewise Fodor (1983, 2000), who otherwise shares few of Evans' assumptions, insists that human thought and concepts are systematic and productive. As a result of thought's component structure, there is no end to the possible combinations of concepts of which humans are capable, and no end to the thoughts that they can thereby think.

It seems, then, that many philosophers are committed to the view that the mind contains a central workspace within which concepts can be freely combined with one another, and in which attitudes of all types can become active, engaging with one another and with systems of inference and decision making. Moreover, since the same philosophers draw a sharp distinction between conceptual thought, on the one hand, and the outputs of our various perceptual modalities, on the other, the concepts and attitudes in question must be *amodal* and non-sensory in character.

Some philosophers go still further, insisting that the central workspace is characterized by a certain kind of rational reflection, and that only concepts and attitudes that are, as one might say, "reflection sensitive" can genuinely count as such. Thus both McDowell (1994) and Brandom (1994, 2000) think that human thought is distinguished by its *spontaneity*, distinctively taking place in a "space of reasons". Our minds (unlike the minds of nonhuman animals, it is said) do not remain passive with respect to our circumstances and experiences. Rather, we are capable of weighing reasons for belief and for action, and any thought of ours can spark an open-ended process of reflection and inference. Indeed, belief, properly so-called, is said to be partly normatively constituted. In order to be capable of believing something, one needs to have some conception of the norms that govern

[1] Fodor (2000) is led to this conclusion because he thinks that cognitive science must traffic in processes that are *local* in character, whereas isotropic processes are, in an important way, *holistic*.

belief-formation. One needs to be capable of reflecting on what one *should* believe in a given evidential context, and of being influenced accordingly.[2]

While Fodor (1998, 2000) would reject the idea that belief is normatively constituted, it seems that philosophers across a very broad spectrum of opinion, who otherwise differ deeply in their methodological commitments, share a conception of the mind as containing, at its core, a workspace in which thoughts can be created, reflected on, and evaluated, and in which attitudes of all types can be active and enter into processes of reasoning and thinking.

I propose to argue that these views are radically mistaken. There is, indeed, a central workspace in the mind whose contents are always conscious. This is so-called "working memory", which has been heavily studied by scientific psychologists for the last 40 years or more. But the operations of working memory are always *sensory* based. Hence with two notable exceptions—namely, sensorily-embedded judgments (seeing *as*, hearing *as*, and so on) and affective feelings directed toward some object or situation—no propositional attitudes figure in the central workspace. Their interactions with other such attitudes are always indirect, mediated by processes that create sensory imagery of various kinds that *can* enter the global workspace. Moreover, with the two exceptions noted above, no propositional attitudes are ever conscious.[3] Indeed, most of the real cognitive work of the mind is done at the unconscious level, where beliefs and goals compete for attentional resources and the control of behavior, and for influence over the contents of the global workspace.

I shall, however, do nothing to challenge the view that cognition is realized in structured amodal representations. On the contrary, I shall assume its truth. So in arguing that working memory is sensory based, I should not be construed as arguing that the *mind* is sensory based, or as defending so-called "sensorimotor" accounts of cognition generally (Barsalou 1999). Rather, the mind contains goals, judgments, decisions, and other attitudes whose realization is "abstract" and amodal. (This is even true of the minds of bees, in my view; see Carruthers 2006, 2009.) And these states can enter into processes of inference and decision making that have limited accessibility. The claim is just that such attitudes can only exert an influence on the global workspace by recruiting suitable forms of sensory imagery.

Developing and providing support for this account will require us to bring together a number of strands of research in cognitive science that are often seen as separate, but are in fact intimately related. One is a "global broadcasting" account of (access) consciousness, initially proposed by Baars (1988, 1997), and since developed and investigated by others. Another is the role of attention directed

2 Compare Peacocke's (1986, 1992) claim that concept possession requires an appreciation of certain inferential moves as being *primitively compelling*. This makes grasp of inferential norms into a condition of genuine concept possession.

3 The notion of consciousness in play here is some or other form of *access*-consciousness (Block 1995). But it matters little for our purposes whether the access in question is characterized in terms of first-order accessibility to processes of belief-formation and decision-making, or instead in terms of higher-order accessibility for immediate self-attribution. This is because in the case of humans, at least, mental states that are "globally broadcast" and hence first-order accessible will at the same time be higher-order accessible, and vice versa.

toward mid-level sensory areas of the brain as the causal instigator of global broadcasting. A third is the well-established finding that perception and imagery in the same modality share the same mechanisms, and can be globally broadcast in the same way. A fourth is the way that concepts can be bound into the contents of globally broadcast perceptual and imagistic states. And fifth, there is the conception of working memory that is now widely accepted by scientists. This regards working memory as a sort of virtual system that uses executive resources located in the frontal and parietal lobes to direct attention toward mid-level sensory areas, issuing in images that can then be sustained, transformed, and manipulated in ways that are globally accessible. These topics will be discussed in turn. We will then confront a number of questions and objections, including the need to show that a sensory-based working memory system is the *only* global workspace that the mind contains.

2 Global broadcasting and attention

Almost all accounts of the nature of phenomenal—or "what-it-is-like"—consciousness can converge on the claim that such consciousness *coincides*, at least, with the global broadcast of perceptual information in the brain, thereby becoming access-conscious. This can be accepted by property dualists and reductive representationalists alike (Tye 1995; Chalmers 1997). Admittedly, some maintain that there can be phenomenally conscious states that are *not* globally broadcast (Block 1995, 2002), but this issue is not really germane to our purposes here. For the form of consciousness that is most closely related to the idea of a global workspace is not phenomenal consciousness, but rather access consciousness.

There is now extensive evidence supporting a global broadcasting account of the conscious accessibility of our perceptual experiences (Baars 1988, 1997, 2002, 2003; Dehaene and Naccache 2001; Dehaene et al. 2001, 2003, 2006; Baars et al. 2003; Kreiman et al. 2003; Sergent et al. 2005; Kouider et al. 2007; Gaillard et al. 2009). Moreover, subsequent analyses of functional connectivity patterns in the human brain have demonstrated just the sort of neural architecture necessary to realize the main elements of a global broadcasting account (Hagmann et al. 2008; Iturria-Medina et al. 2008; Bullmore and Sporns 2009; Gong et al. 2009; Shanahan 2010). Specifically, these studies show the existence of a long-range "connective core" along the midline of the brain, which serves to link a set of more densely locally connected brain areas. The connective core facilitates the widespread disbursal of sensory information, while at the same time serving as a bottleneck, forcing the local areas to compete with one another to have their messages transmitted through it.

Among the experiments that support a global broadcasting account are ones that employ carefully matched masked and unmasked stimuli, resulting in unconscious and conscious perception respectively (Dehaene et al. 2006). In the masked condition there is reverberating activity in areas of occipital, posterior parietal, and temporal cortex. This results in significant priming effects, suggesting that the stimuli have been processed up to the semantic level. Subjects nevertheless remain unaware of having seen anything, and no relevant activity in wider areas of the cortex occurs. In the unmasked condition, in contrast, the activity in visual cortex is swiftly followed by widespread patterns of reverberation in temporal, frontal, and parietal lobes. The result is that subjects are aware of having seen the stimulus, and can report on its properties.

The same body of research cited above also suggests that the main determinant of global broadcasting is attention. (See Awh et al. 2006; Knudsen 2007 for reviews.) This can be bottom–up, caused by highly salient stimuli, sudden changes in the environment, or stimuli of innate or learned emotional significance. (Examples include the proverbial snake in the grass, a loud noise, a smiling face, or the sound of one's own name.) Or it can be top-down, driven by the subject's high-level goals and interests in the circumstances. Moreover, attention seems to be specific in the loci of its focus, which is on mid-level sensory processing areas. In the case of vision, this would include the regions that process color, form, motion, spatial layouts, and faces, but not the primary visual projection area V1, and not the regions of temporal and parietal association cortices that receive output from mid-level regions. This point will prove to be of some importance later.

In addition, we are even beginning to understand *how* attention can initiate the global broadcasting of sensory information. It seems that attention modulates the activity of populations of neurons that code for the attended-to properties. Thus presentation of a spatial cue while subjects are remembering information from a number of locations reduces electrical activity in posterior parietal visual cortex (since fewer items now need to be recalled) while dramatically improving conscious recall of the target items (Kuo et al. 2012). And likewise, presentation of a category cue while subjects are remembering both faces and scenes modulates activity in the fusiform face area or in the parahippocampal visual scene area respectively (Lepsien et al. 2011). The mechanism underlying these findings seems to be that top–down attention directed at mid-level sensory areas serves to increase the stimulus specificity of populations of neurons that are tuned to task-relevant stimuli, while damping down the responses of neural populations that respond to aspects of the stimulus that are not task-relevant, and while boosting the overall activity of the former populations across the critical threshold for global broadcasting (Gazzaley et al. 2005; Knudsen 2007).

Moreover, the contributions made by the different components of the attentional network are also increasingly well understood. First, there appear to be *two* such networks (albeit interacting with one another), for top–down and bottom–up forms of attention respectively (Corbetta and Shulman 2002). The top–down network links prefrontal cortex (and especially the frontal eye fields in the case of vision) with the intraparietal sulcus bilaterally. It seems that the contribution of prefrontal cortex is to maintain the goals and expectations of the organism, transmitting to parietal cortex attentional priorities and representations of attentional relevance. It also seems that regions around the top of the intraparietal sulcus are critically involved in the manipulative aspects of working memory and are highly interconnected with prefrontal cortex (Postle et al. 2006; Koenigs et al. 2009), whereas a closely connected region lower on the intraparietal sulcus modulates activity in visual cortex (Uddin et al. 2010). The bottom–up network, in contrast, is centered on the right temporo-parietal junction, interacting with areas of ventral prefrontal cortex, as well as with the parietal areas of the top–down network (Corbetta and Shulman 2002).

3 Percepts, concepts, and imagery

Although attention directed at mid-level perceptual areas is the main causal determinant of global broadcasting, this is not to say that the contents of global broadcasts are restricted to the outputs of such areas. On the contrary, there is every reason to think that conceptual information that is activated by interactions between mid-level areas and the association areas (especially temporal cortex in the case of vision) gets bound into the content of attended perceptual states and is broadcast along with the latter. Hence we don't just see a spherical object moving along a surface, but a tomato rolling toward the edge of the counter top; and we don't just hear a sequence of phonemes when someone speaks, but we hear what they are saying; and so on.

Kosslyn (1994) characterizes high-level visual processing as a kind of back-and-forth "questioning" of mid-level visual information by conceptual systems charged with classification of the properties of the input. As each successful classification is made, it gets bound into the content of the visual percept in question, and is thereby made available for classification by yet other conceptual systems. The result can be a cascade of increasingly abstract classifications, with each property, once identified, being added to an indexical "object file" for the object or event in question. All of this takes place unconsciously, and the results may fail to be globally broadcast. But when attention is appropriately directed, they can be globally broadcast to multiple systems, thereby becoming access-conscious. The result is that what we see is an amalgam of conceptual and nonconceptual information—we see the fine-grained colors and textures in a rose petal, for example, while also seeing it *as* a rose petal.

Perception of speech is known to have very similar properties (Hickok and Poeppel 2007). Nonconceptual representations of sound begin to interact quite early in the auditory system with structures in the language faculty. The latter parses the sound stream into phonemes and words, and assigns syntactic and semantic representations to it. In addition, the emerging representations interact with aspects of the mindreading faculty, attempting to figure out the speaker's intentions, such as whether the utterance is ironic or literal, and so on. And probably all of these processes take place in parallel (or at least in complex feedback loops) rather than in a linear order from phonemes, to lexical items, to syntax, to semantics, to speaker meaning. The upshot, as in the case of vision, is that by the time the auditory representations are globally broadcast they have (normally) already been conceptualized and interpreted. Hence we hear the speaker *as* saying one thing rather than another. The heard content of the speech is bound into the representation of the sound stream. Indeed, the former will normally be the central focus of our awareness, with the particular phonemes employed dropping out to the periphery.

Another important finding for our purposes is that imagery in a given sense modality re-uses the mid-level processing areas of that modality (Paulescu et al.

1993; Kosslyn 1994; Shergill et al. 2002; Kosslyn et al. 2006). There seem to be two main ways in which imagery can be generated. One is by using conceptual templates to direct attention at the set of micro-regions that would normally perceptually encode the object or event represented. When stimulated by attention, these regions become sufficiently active to be globally broadcast, thereby making the information that they encode consciously accessible. The other major way in which imagery is created is through off-line activation of motor schemata. Efference copies of a set of motor instructions are initially generated (but with output to motor systems suppressed), being used to create a sensory "forward model" of the likely consequences of the movement (Jeannerod 2006). When attended to, this representation, too, can be globally broadcast. Such imagery can be proprioceptive, as when one imagines how it would feel to execute a particular movement, or it can be visual or auditory. Especially important in the latter case are the auditory images that result from off-line activation of instructions for producing speech, which result in auditory representations of the speech act that would normally result, in so-called "inner speech" (Paulescu et al. 1993; Shergill et al. 2002).

4 Working memory

Note that the various findings outlined in Sects. 2 and 3 have been introduced while hardly mentioning working memory. Nevertheless, they converge to help make sense of, and to lend independent support for, an account of working memory that is now widely accepted. This is that working memory depends on frontal-lobe attentional and executive systems utilizing top–down signals to recruit activity in mid-level perceptual regions of the brain, thereby creating endogenously-generated activity within a global workspace that is accessible to all of the various systems that would normally consume and respond to perceptual input (Müller and Knight 2006; Postle 2006; D'Esposito 2007; Knudsen 2007; Jonides et al. 2008; Sreenivasan et al. 2011).

In Baddeley's classic model, working memory consisted of a central executive deploying two "slave" systems—the phonological loop and the visuo-spatial sketchpad (Baddeley and Hitch 1974; Baddeley 1986). These systems were thought of as specialized subsystems of the central executive, and were believed to be located in or close to the frontal lobes that house the executive. But this initial model has come under pressure in three ways. One derives from Baddeley himself. He has since proposed the existence of an "episodic buffer" that serves to integrate episodic and semantic conceptual information with the contents of the phonological loop and visuo-spatial sketchpad (Baddeley 2006). This seems to mirror the view from the vision and language sciences, that conceptual information can be bound into perceptual states and broadcast along with them.

A second challenge to Baddeley's early model is that evidence has been increasingly accumulating that the "slave" systems of the central executive are none other than the auditory and visual systems, whose resources can be recruited by executive activity to broadcast and sustain imagistic representations of the relevant sort. And then the third critique is that it is now believed that working memory can deploy the resources of any sensory system, not just vision and audition. Hence there is smell-based working memory, tactile working memory, proprioceptive working memory, affective working memory, and so on (Dade et al. 2001; Harris et al. 2002; Jeannerod 2006; Mikels et al. 2008).

The generally accepted current picture, then, is of a set of executive systems that deploy attentional and other resources to recruit activity in mid-level sensory areas of the brain, resulting in globally broadcast representations that can be sustained, manipulated, or replaced by further actions of the executive. If this account is correct then there is, indeed, a central workspace of the mind in which information can be made widely accessible, and within which representations can be endogenously generated and manipulated in the service of the organism's goals. But it is a sensory-based system. What figures within it are not (in general) propositional attitudes, but rather visual images, auditory images, imagined movements, and so forth. There are, however, some exceptions to this claim, which we discuss next.

4.1 Attitudes in working memory

There are two exceptions to the generalization that no attitudes figure directly in working memory. One is that conceptual information bound into the content of a perceptual or imagistic state can on some views qualify as a perceptual *judgment*. When one sees a moving object in one's kitchen *as* a rolling tomato, for example, this is apt to give rise to a stored belief-state with the content, *a tomato was rolling across the counter top*, and it is likewise available to guide planning and action-selection in the manner of a judgment. Similarly, when conceptual contents are bound into the representation of heard speech one might hear someone *as* asking the way to the church, for instance (hearing this as a question with a specific content, resulting from the interpretive work of the language and mindreading faculties). This can perhaps qualify as a sensorily-embedded judgment with the content, *he is asking the way to the church*. Similar points will then obtain with respect to the deployment of visual imagery or inner speech in the contents of working memory.

The second exception to the generalization that working memory does not traffic in propositional attitudes is that in addition to sensorily-embedded judgments it can contain momentary desires. For *affect*, too, can be globally broadcast and hence become widely accessible (Mikels et al. 2008). One dimension of affect is *arousal*. One can be aware, not just of the bear that one sees looming out of the bushes, but also some of the physiological reactions that are occasioned by one's fear. The other main dimension of affect is *valence*, which is best thought of as a form of nonconceptual representation of value (Carruthers 2011). Hence the presence of the bear will *seem bad* as a result of the negative valance that is broadcast alongside of one's perception of it.[4]

[4] Notice that I do not say that valence is *bound into* the contents of perception. This is because the evidence suggests that it is not. On the contrary, valence that is a product of many different sources (one's background mood, irrelevant features of the stimulus or its surroundings, and so on) is by default taken to be directed at whatever is the current object of attention (Schwarz and Clore 1983, 2003; Forgas 1995; Higgins 1997; Gasper and Clore 2000; Winkielman et al. 2005; Li et al. 2007; Schnall et al. 2008). But the result is a motivational state that functions somewhat like an active desire as philosophers traditionally conceive of it, except that the relationship between one's experienced affect and one's stable values is highly labile and context dependent.

With these two exceptions noted, it nevertheless remains that case that working memory does not contain any judgments or beliefs that are *not* sensorily embedded. One can *experience* oneself as judging that the church is a mile to the south when one entertains an auditory image of the sentence, "The church is a mile south", just as one can experience someone as asking the way to the church when one hears him say, "Which way is the church?" But this imagined sentence is not itself a judgment (any more than another person's speech is itself a judgment). At best it expresses such a judgment, which is a distinct token event from the item of inner speech in question. Nor does working memory ever contain any decisions, intentions, values, or goals. So the philosophical conception of central cognition outlined in Sect. 1 is directly challenged by this account.

4.2 Attitude contents in working memory

It might be felt that episodic memories should be added to the list of attitudes that can figure actively in working memory. For when one entertains an episodic memory one *re-lives*, to some significant degree, an earlier experience. Such memories consist partly of stored images that recapitulate some of the experiences one had at the time. If this is so, then it might seem that episodic memories, themselves, can be activated into working memory. In effect, such memories would be sensorily-embedded episodic memory *judgments*.

This interpretation should be resisted, however. Granted, the *content* of an episodic memory image can derive from a stored representation of an earlier experience, at least in general, and to some significant degree.[5] But it is another matter to claim that such episodes automatically have the causal role distinctive of episodic memory (grounding further inferences about the past and guiding actions that depend upon the past). Indeed, the evidence suggests that what constitutes an episodic image as a *memory* image is mindreading-based interpretation. For the kinds of representation that are activated in episodic memory retrieval are also active in episodes of imagination and fantasy. Yet the resulting images don't wear their provenance on their sleeves.

The extensive empirical literatures on source monitoring (Kunda 1999; Mitchell and Johnson 2000) and on metacognition (Dunlosky and Metcalfe 2009) suggest that the products of imagination and of episodic memory can only be distinguished by a variety of sensory cues that are internal to those images themselves (such as greater specificity and vividness) or that tend to accompany memories (such as feelings of familiarity), together with aspects of the surrounding context. Hence the images in question can only acquire the role *of* an active memory or *of* an imagined situation by being interpreted as such. As a result, what figures in working memory is not itself an episodic memory, but rather a sensory image that is caused by such a memory and has been interpreted as such.

[5] The qualifications are needed because of the well-known constructive nature of memory. This means that any memory image will consist at least partly in elements that have been added later, or that have been added during the process of activating and recovering the memory itself. See Schacter (2001).

These claims are fully consistent with the finding that episodic remembering shares the same cognitive and neural resources as *prospection* (the imagining of future scenarios; Gilbert and Wilson 2007; Schacter et al. 2007, 2008). We now know that future-directed fantasy and imagination draw on the same hippocamal and para-hippocampal brain structures that underlie episodic memory (Buckner 2010). Indeed, some have claimed that the properties of the episodic memory system can only be understood by seeing how it has been shaped in the service of effective prospection and future planning (Schacter et al. 2008). So it makes good sense that there should be no intrinsic differences between episodic and prospective images, and that what gives an image the one sort of functional role rather than the other is a result of interpretation.

It seems that while the *content* of an episodic memory can be directly active in working memory, the attitudinal component (which constitutes it *as* a memory) is not. Something similar is true of semantic memories whose content can be stored in, or directly bound into, a visual format. (We will return to consider the relationship between belief and inner speech in Sect. 5.2, which raises different considerations.) Indeed, it seems that anything one can *see as* being the case one can also *remember as* being the case. If one can see a particular house *as* Sally's home, for example, then a visual image with the same content figuring in working memory can be an activation of the content of one's *belief* about where Sally lives. (But again, what gives the image a belief-like rather than a supposition-like causal role will depend on one's interpretation of its nature.)

5 Objections and elaborations

A number of questions and concerns can be raised about the account of working memory sketched in Sect. 4. The main ones are whether sensory activity merely *accompanies* working memory (rather than being causally necessary for its operation) and whether sensory-based working memory is the *only* global workspace that the human mind contains. These and related issues will now be discussed.

5.1 Objection 1

Someone might grant that working memory always *implicates* sensory representations of one sort or another, while claiming that these representations are not necessary for the system to function. On this account, working memory could contain amodal attitude events of various sorts (goals, decisions, non-sensory judgments, and so on), just as philosophical orthodoxy conceives. But these events would serve to spread activation into sensory areas in a content-related manner, so that working memory activities would always be accompanied by sensory images. This would explain why brain imaging studies always show activity in sensory areas during working memory tasks while undermining the claim that working memory is sensory *based*.

The idea of spreading activation may indeed be well motivated in other contexts. For example, Zwaan et al. (2002) show that subjects presented with a sentence such as, "He saw an eagle in the sky" are faster to respond to a picture of an eagle with its wings outstretched than they are if presented with a sentence such as, "He saw an eagle in the tree." This suggests that a visual representation of a flying eagle had already been activated (whether consciously or unconsciously). Similarly, Mahon and Caramazza (2008) review a variety of kinds of evidence suggesting that activation of tool concepts always covaries with activation of the motor plans for using those tools, as well as visual representations of the tools themselves and their standard uses. They nevertheless argue that there is an abstract, amodal, level of conceptual representation involved, with activation spreading between this and associated sensory and motor representations. They show, for example, that there are patients who cannot recognize tools (sensory deficit) who can nevertheless both say and pantomime how a tool of a given sort should be used. At the same time other patients cannot use tools (motor deficit), while nevertheless being able to recognize both tools and pantomimes of their standard uses.

Although the idea of spreading activation makes sense, the present objection nevertheless runs up against the finding that attention directed at mid-level sensory areas is a necessary (and perhaps sufficient) condition for global broadcasting to occur, and hence for entry of a given conceptual representation into working memory (Awh et al. 2006; Knudsen 2007; Prinz 2012). If this is so, then sensory activation is by no means ancillary to the operations of working memory, but rather constitutes its very foundation.

A number of findings corroborate this claim. One is that in macaque monkeys the main brain areas involved in the regulation of top–down control of visual attention (especially the frontal eye fields) project *only* to mid-level visual cortical areas (Stanton et al. 1995). They do not project to regions of temporal cortex that would most likely realize visually-based judgments, suggesting that such judgments need to be bound into attended sensory information in order to become globally accessible. This conclusion is not forced upon us, of course. For it may be that some of the other components of the top–down attentional system project to non-sensory regions. Or more radically still, it may be that one of the phylogenetic changes to have occurred since the last common ancestor of ourselves and macaques was the evolution of neural systems to support the global broadcast of conceptual information directly, independent of sensory attention.

In fact in humans, however, the top–down attentional system consists of a network of prefrontal cortical areas (including the frontal eye fields in the case of vision) together with regions around the intraparietal sulcus (Corbetta and Shulman 2002; Todd and Marois 2004; Knudsen 2007; Egner et al. 2008), as we noted in Sect. 2. These are strongly interconnected with one another. And the latter is strongly connected to mid-level visual areas but *not* to temporal cortex (Uddin et al. 2010). This again suggests that attention directed toward visual cortex (or toward mid-level sensory areas in other sensory modalities; Pasternak and Greenlee 2005) is a necessary condition for entry of information into working memory.

Moreover, a variety of kinds of evidence suggest that damage or disruption to mid-level sensory areas has a corresponding impact on working memory (Pasternak and Greenlee 2005). For example, brain damage to these areas impairs working memory abilities in the corresponding domain (Levine et al. 1985; Gathercole 1994; Müller and Knight 2006). Moreover, transcranial magnetic stimulation (TMS) applied to mid-level sensory areas disrupts working memory tasks that would normally involve imagery of the relevant sort. Thus TMS applied to posterior parietal cortex during the retention interval disrupts working memory for spatial layouts (Oliveri et al. 2001), whereas TMS applied to area V5/MT disrupts working memory for direction of motion (Campana et al. 2002). In addition, presentation of distracter stimuli during the retention interval in working memory tasks only has a disruptive effect on those tasks if the stimuli are perceptually similar to the target. For example, location distracters interfere with memory for location but not color, whereas color distracters interfere with working memory for color but not location (Vuontela et al. 1999).

It might be objected that such results are not really surprising, given that the working memory tasks in question themselves involve sensory properties. Hence it would be consistent with such results that working memory of a semantic sort, or involving abstract properties, should *not* depend on the activity of mid-level sensory areas. But this objection seems not to be sustainable. For notice that many of the results reported above involve spatial working memory (e.g. concerning direction or position). Yet space is an abstract property that can be represented in many different sense modalities. So if there were such a thing as a non-sensory form of working memory, one would expect that spatial working memory would be a good candidate. But as we have seen, the experimental results seem to rule this out.

It would appear, then, that the activity that regularly shows up in sensory areas during imaging studies of working memory is by no means merely a result of spreading activation from the presence of globally broadcast amodal propositional attitudes. It is, rather, a necessary condition for amodal, abstract, or concept-involving representations to be entertained in working memory at all.

5.2 Objection 2

It might be argued that the evidence that sensory representations are causally indispensable to working memory can be accommodated, consistent with claiming that working memory is a space in which propositional attitudes of all sorts can be tokened and can interact. For it might be said that attended sensory activity is merely a necessary condition for the attitudes in question to become globally broadcast. On this account a non-sensory judgment or decision would need to recruit sensory activity of some suitable sort into which it can be bound in order to gain access to the workspace. Thus a judgment that the church is due south might issue in the imagined sentence, “The church is due south.” But the token judgment that causes the production of that item of inner speech is *itself* globally broadcast as a result.

One problem for such an account is that it is hard to understand how a represented sentence can have a judgment with the same content bound into it. (The fact that the sentence can be heard as *expressing* a judgment that the church is due south is another matter. This was acknowledged earlier. What is problematic is to see how the sentence can be—or can contain—the *judgment* that the church is due south.) For consider overt utterances. When someone says aloud that the church is due south we are not tempted to think that this performance *is* a judgment, nor that it *contains* one. Rather, at best it *expresses* a judgment which occurred slightly earlier and which caused the utterance in question. Likewise, then, for subvocal or inner speech.

In addition, there is evidence that language *comprehension* as well as language production areas of the brain are active in inner speech (Paulescu et al. 1993; Shergill et al. 2002). While the represented sentence is a forward model of the likely result of a particular speech act (which is itself suppressed), that sensory representation seemingly needs to be received by the language comprehension system and interpreted. This process is the same in principle as that involved when interpreting the speech of another (albeit with differences of detail). Admittedly, the fact that a particular judgment was active just prior to the represented utterance is likely to greatly aid the process of interpretation. For the latter is driven in large part by the relative *accessibility* of lexical, semantic, and syntactic properties (Sperber and Wilson 1995). Hence the language comprehension process in inner speech will be heavily biased by the immediately preceding production process. But the token judgment (if there was one) that originally served as a partial cause of the production process is in no sense preserved and bound into the result (Carruthers 2011).

Moreover, if non-sensory judgments, decisions, and so forth were frequently tokened in working memory, as standard philosophical models assume, then reporting the occurrence of those attitude events should be a straightforward and reliable matter. For those attitudes would be globally broadcast in a way that makes them available to whatever systems engage in self-attribution of mental states and verbal report. But this is not what we find. On the contrary, the literature in social psychology is rife with experiments in which people make mistaken reports of their current or very recently past thoughts, in circumstances where those thoughts should surely have been active in working memory, if such a thing were ever possible at all. Carruthers (2011) argues at length that this pattern of findings is best explained by an account that does *not* postulate the global broadcasting of propositional attitude events in working memory.

5.3 Objection 3

Someone might concede these points while claiming that the resulting interpreted sentences can acquire a causal role appropriate for a judgment, or a decision, *subsequent* to the process of interpretation. By hearing myself as expressing a judgment or decision, I can make it the case that I *am* making that judgment or decision. I can do this by constraining my own future behavior (whether inner or outer) in appropriate ways. This possibility is explored and criticized extensively in Carruthers (2011), and will be discussed only briefly here.

In short, the reason why these interpreted imagistic events don't qualify as judgments or decisions in their own right is that they lack the right sort of functional profile. For the best account that we have of *how* interpreted images achieve their attitude-like effects is that the latter result from the intervention of higher-order beliefs and goals (which remain unconscious; Frankish 2004, 2009). Suppose, for example, that the sentence, "I shall go to the bank", is tokened in inner speech. Under interpretation by the language faculty working together with the mindreading system, this might be heard as expressing a *decision* to go to the bank. This is, note, a *higher-order* representation: I *represent* myself *as* making a decision. The event certainly isn't a decision of the regular sort.

How, then, might such an event issue in bank-going behavior? Well, if I have a standing desire to do what I have decided, or if I am motivated to be a decisive or strong-willed person, then the belief that I have made a decision may evoke a higher-order motive to execute that decision. These taken together might then cause me to walk over to the bank, or to form the intention of visiting the bank. Notice, though, that the underlying practical reasoning will look something like this: I have decided to go to the bank; I want to do what I have decided; so I shall go to the bank. The final event in this sequence is *itself* a decision to go to the bank. This makes it quite unlikely that the interpreted item of inner speech that initiated the sequence should *also* be an instance of the very same type of attitude—namely, a decision to go to the bank.

Carruthers (2011) argues that one of our tacit commitments regarding the nature of decisions is that they should issue in actions or intentions directly, without being mediated by further reasoning about whether or not to act. On the contrary, we think that a decision is what concludes an episode of practical reasoning and *settles* whether one will act (Bratman 1987, 1999). But in the case of an item of inner speech that is heard (under interpretation) as expressing a decision, this only leads to action or to the formation of an intention in the indirect manner outlined above. That event, by itself, does not settle anything. Carruthers (2011) argues that this result generalizes, in such a way that none of the representations that figure in sensory-based working memory (besides sensorily-embedded judgments and affective feelings) qualifies as a kind of propositional attitude.

5.4 Objection 4

Someone might concede the existence of a sensory-based working memory system while claiming that we *also* possess an attitudinal working memory system that can operate independently of sensory representations. But this seems like an ad hoc attempt to shore up the traditional philosophical picture. For we have no evidence of the existence of such a system. Although human working memory has been extensively investigated over the last half century, no evidence of a purely attitudinal system has emerged, and no theorists in psychology (as opposed to artificial intelligence; see Sect. 5.5) have proposed such a thing.

Moreover, what working memory is fundamentally *for* is general-purpose problem solving, of just the sort that is standardly measured by tests of fluid general intelligence, or *g*. So if there were an attitudinal working memory system in addition to a sensory-based one, we would expect that variance in the properties of that system across individuals should account for some of the variance in general intelligence. But it seems that this is not the case. For correlations between sensory-based working memory ability (measured using the sorts of sensory-involving tasks that are used to explore its nature) and fluid *g* are generally quite high. Indeed, in a very large study involving hundreds of participants, Colom et al. (2004) found that variance in the one could account for no less than 96 % of the variance in the other. This suggests quite strongly that there is just a single working memory system, which is the sensory-based one that has been explored by psychologists.[6]

5.5 Objection 5

Our final objection takes the form of a puzzle: *why* should working memory be sensory based? This seems inexplicable from a reverse-engineering standpoint. If one wanted to design a system that could reason about abstract matters and integrate the results so as to issue in decisions that are in turn modulated by prior intentions, then there seems no reason why one should give it a sensory basis. On the contrary, the system would most naturally be designed using only abstract, amodal, representations—a language of thought, no less. And indeed, researchers in artificial intelligence who attempt to develop flexible reasoning and decision making systems use just such a strategy (Newell 1990; Anderson 1993; Lovett et al. 1999; Young and Lewis 1999).

The solution to the puzzle, however, is an evolutionary one, and depends upon the principle that the design of a system can be constrained by ancestral architectures. In the present instance, systems that support the global broadcast of attended sensory information are a long-standing feature of mammalian (and perhaps avian) brains. These were then co-opted to constitute a simple form of short-term memory, within which recently-experienced representations could be attentionally sustained in the absence of a current stimulus, as when a monkey watches an object roll behind an occluder (Santos et al. 2005). This provided the foundation on which working memory could evolve, issuing first in simple forms of endogenously generated imagery, and developing increasingly sophisticated abilities to transform and manipulate that imagery (Coolidge and Wynn 2009). What explains the fact that the human working memory system is sensory based is that it evolved from simpler forms of sensory short-term memory.

I have argued elsewhere that the minds of non-human animals (not only birds and mammals, but also invertebrates) contain a variety of propositional attitudes realized in structured amodal representations, together with limited forms of inference and decision making defined over those representations (Carruthers 2006, 2009). The difference between human and animal minds is not that we can think while they cannot; nor is it that our attitudes are inferentially promiscuous while theirs are not. Rather, the major difference that is relevant to our present discussion

[6] Colom et al. (2010) reason that if general intelligence is a direct function of the efficiency of sensory-based working memory, then training in the latter should improve the former. However, they failed to find such effects. But this may be because they did not separately analyze those who benefited from the working-memory training and those who did not. For Jaeggi et al. (2011) find robust improvements in fluid general intelligence over a 3 month interval among children who show the greatest benefit from training in working memory tasks.

is best seen as an evolutionary kludge: humans (and perhaps other primates) have evolved a sensory-based working memory system, providing a workspace through which all the various components of the mind can interact with one another (Carruthers 2013). It may not be an elegant solution to the comparative inflexibility of animal minds, but it works.

6 Conclusion

I have suggested that philosophical assumptions about central cognitive processes are radically mistaken. There is no workspace within which attitudes of all kinds can be active and interact with one another. Nor are most kinds of attitude inferentially promiscuous, if this is interpreted to mean that those attitudes themselves (rather than their sensory effects) must be capable of entering into inferences with all other attitudes. In addition, if conscious attitudes are thought to be those that are globally accessible, then there are at most two kinds of conscious attitude: perceptually-embedded judgments and affective feelings directed at some object or event. All other forms of propositional attitude are always unconscious (although in some cases their *contents* can be conscious, as with episodic and some kinds of semantic memory).

Although our attitudes are not isotropic in the sense that Fodor (1983) intends, the distinctive flexibility of human cognitive processes can arguably be explained, nevertheless. For although most kinds of attitude cannot interact with others directly, they can still have an influence on the contents of the sensory-based global workspace. For example, a belief or non-sensory judgment might issue in an episode of inner speech with the same or sufficiently similar content. When interpreted, the latter can assume a causal role somewhat like that of a judgment (Frankish 2004), and its content will be accessible to all of the systems that consume global broadcasts. This might activate additional stored information or relevant goals, issuing in yet further episodes of inner speech or in the evocation of suitable visual or other imagery. As a result, it is true that any belief or other attitude can be brought to bear in the evaluation of any belief or decision *indirectly*, through the effects that attitudes can have on the contents of the global workspace. Carruthers (2006) argues at length that the distinctive flexibility of human cognition can be explained adequately in such terms, thereby undercutting any need to postulate a non-sensory central workspace (while also undercutting Fodor's principled pessimism about the prospects for a cognitive science of central cognition).

As already noted, it also follows from the account defended here that most kinds of attitude remain entirely below the surface of our conscious lives, engaging in limited interactions with others, but helping to shape or motivate the global broadcast of sensory-involving images of various sorts. As such, those attitudes are never actively controlled nor directly normatively constrained. Any philosophical positions that are premised on such views will therefore need to be re-examined.

Acknowledgments Some of the material used in this article is drawn from Carruthers (2011). I am grateful to Lizzie Schechter and Wayne Wu for their comments on an earlier draft.

References

Anderson, J. (1993). *Rules of the mind*. Hillsdale: Erlbaum.

Awh, E., Vogel, E., & Oh, S. (2006). Interactions between attention and working memory. *Neuroscience, 139*, 201–208.

Baars, B. (1988). *A cognitive theory of consciousness*. New York: Cambridge University Press.

Baars, B. (1997). *In the theatre of consciousness*. New York: Oxford University Press.

Baars, B. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Sciences, 6*, 47–52.

Baars, B. (2003). How brain reveals mind: Neuroimaging supports the central role of conscious experience. *Journal of Consciousness Studies, 10*, 100–114.

Baars, B., Ramsoy, T., & Laureys, S. (2003). Brain, consciousness, and the observing self. *Trends in Neurosciences, 26*, 671–675.

Baddeley, A. (1986). *Working memory*. Oxford: Oxford University Press.

Baddeley, A. (2006). *Working memory, thought, and action*. Oxford: Oxford University Press.

Baddeley, A., & Hitch, G. (1974). Working memory. In G. Bower (Ed.), *Recent advances in learning and motivation, vol. 8* (Vol. 8). New York: Academic Press.

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

Block, N. (1995). A confusion about the function of consciousness. *Behavioral and Brain Sciences, 18*, 227–247.

Block, N. (2002). The harder problem of consciousness. *The Journal of Philosophy, 99*, 1–35.

Brandom, R. (1994). *Making it explicit*. Cambridge: Harvard University Press.

Brandom, R. (2000). *Articulating reasons*. Cambridge: Harvard University Press.

Bratman, M. (1987). *Intentions, plans, and practical reason*. Cambridge: Harvard University Press.

Bratman, M. (1999). *Faces of intention*. Cambridge: Cambridge University Press.

Brewer, B. (1999). *Perception and reason*. Oxford: Oxford University Press.

Buckner, R. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of Psychology, 61*, 27–48.

Bullmore, E., & Sporns, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience, 10*, 186–198.

Camp, E. (2004). The generality constraint, nonsense, and categorical restrictions. *Philosophical Quarterly, 54*, 209–231.

Campana, G., Cowey, A., & Walsh, V. (2002). Priming of motion direction and area V5/MT: A test of perceptual memory. *Cerebral Cortex, 12*, 663–669.

Carruthers, P. (2006). *The architecture of the mind*. Oxford: Oxford University Press.

Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz (Ed.), *The philosophy of animal minds*. Cambridge: Cambridge University Press.

Carruthers, P. (2011). *The opacity of mind*. Oxford: Oxford University Press.

Carruthers, P. (2013). Evolution of working memory. *Proceedings of the National Academy of Sciences, 110*, 10371–10378.

Chalmers, D. (1997). Availability: The cognitive basis of experience. *Behavioral and Brain Sciences, 20*, 148–149.

Colom, R., Quiroga, M., Shih, P., Martínez, K., Burgaleta, M., Martínez-Molina, A., et al. (2010). Improvement in working memory is not related to increased intelligence scores. *Intelligence, 38*, 497–505.

Colom, R., Rebollo, I., Palacios, A., Juan-Espinosa, M., & Kyllonen, P. (2004). Working memory is (almost) perfectly predicted by g. *Intelligence, 32*, 277–296.

Coolidge, F., & Wynn, T. (2009). *The rise of Homo sapiens: The evolution of modern thinking*. Maiden: Wiley-Blackwell.

Corbetta, M., & Shulman, G. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3*, 201–215.

D'Esposito, M. (2007). From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society B, 362*, 761–772.

Dade, L., Zatorre, R., Evans, A., & Jones-Gottman, M. (2001). Working memory in another dimension: Functional imaging of human olfactory working memory. *Neuroimage, 14*, 650–660.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences, 10*, 204–211.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*, 1–37.

Dehaene, S., Naccache, L., Cohen, L., Bihan, D., Mangin, J., Poline, J., et al. (2001). Cerebral mechanisms of word priming and unconscious repetition masking. *Nature Neuroscience, 4*, 752–758.

Dehaene, S., Sergent, C., & Changeux, J. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences, 100*, 8520–8525.

Dunlosky, J., & Metcalfe, J. (2009). *Metacognition*. Thousand Oaks: Sage Publications.

Egner, T., Monti, J., Trittschuh, E., Wieneke, C., Hirsch, J., & Mesulam, M. (2008). Neural integration of top-down spatial and feature-based information in visual search. *Journal of Neuroscience, 28*, 6141–6151.

Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.

Fodor, J. (1983). *The modularity of mind*. Cambridge: MIT Press.

Fodor, J. (1998). *Concepts*. Oxford: Oxford University Press.

Fodor, J. (2000). *The mind doesn't work that way*. Cambridge: MIT Press.

Forgas, J. (1995). Mood and judgment. *Psychological Bulletin, 117*, 39–66.

Frankish, K. (2004). *Mind and supermind*. Cambridge: Cambridge University Press.

Frankish, K. (2009). Systems and levels. In J. Evans & K. Frankish (Eds.), *In two minds*. Oxford: Oxford University Press.

Gaillard, R., Dehaene, S., Adam, C., Clémenceau, S., Hasboun, D., Baulac, M., et al. (2009). Converging intracranial markers of conscious access. *PLoS Biology, 7*, 472–492.

Gasper, K., & Clore, G. (2000). Do you have to pay attention to your feelings to be influenced by them? *Personality and Social Psychology Bulletin, 26*, 698–711.

Gathercole, S. (1994). Neuropsychology and working memory: A review. *Neuropsychology, 8*, 494–505.

Gazzaley, A., Cooney, J., McEvoy, K., Knight, R., & D'Esposito, M. (2005). Top-down enhancement and suppression of the magnitude and speed of neural activity. *Journal of Cognitive Neuroscience, 17*, 507–517.

Gilbert, D., & Wilson, T. (2007). Prospection: Experiencing the future. *Science, 317*, 1351–1354.

Gong, G., He, Y., Concha, L., Lebel, C., Gross, D., Evans, A., et al. (2009). Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography. *Cerebral Cortex, 19*, 524–536.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, C., et al. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology, 6*(7), e159.

Harris, J., Miniussi, C., Harris, I., & Diamond, M. (2002). Transient storage of a tactile memory trace in primary somatosensory cortex. *Journal of Neuroscience, 22*, 8720–8725.

Hickok, G., & Poeppel, D. (2007). The cortical organization of speech processing. *Nature Reviews Neuroscience, 8*, 393–402.

Higgins, E. (1997). Beyond pleasure and pain. *American Psychologist, 52*, 1280–1300.

Hurley, S. (2006). Making sense of animals. In S. Hurley & M. Nudds (Eds.), *Rational animals?* Oxford: Oxford University Press.

Iturria-Medina, Y., Sotero, R., Canales-Rodriguez, E., Aleman-Gomez, Y., & Melie-Garcia, L. (2008). Studying the human brain anatomical network via diffusion-weighted MRI and graph theory. *NeuroImage, 40*, 1064–1076.

Jaeggi, S., Buschkuehl, M., Jonides, J., & Shah, P. (2011). Short- and long-term benefits of cognitive training. *Proceedings of the National Academy of Sciences, 108*, 10081–10086.

Jeannerod, M. (2006). *Motor cognition*. Oxford: Oxford University Press.

Jonides, J., Lewis, R., Nee, D., Lustig, C., Berman, M., & Moore, K. (2008). The mind and brain of short-term memory. *Annual Review of Psychology, 59*, 193–224.

Knudsen, E. (2007). Fundamental components of attention. *Annual Review of Neuroscience, 30*, 57–78.

Koenigs, M., Barbey, A., Postle, B., & Grafman, J. (2009). Superior parietal cortex is critical for the manipulation of information in working memory. *Journal of Neuroscience, 29*, 14980–14986.

Kosslyn, S. (1994). *Image and brain*. Cambridge: MIT Press.

Kosslyn, S., Thompson, W., & Ganis, G. (2006). *The case for mental imagery*. New York: Oxford University Press.

Kouider, S., Dehaene, S., Jobert, A., & Le Bihan, D. (2007). Cerebral bases of subliminal and supraliminal priming during reading. *Cerebral Cortex, 17*, 2019–2029.

Kreiman, G., Fried, I., & Koch, C. (2003). Single neuron correlates of subjective vision in the human medial temporal lobe. *Proceedings of the National Academy of Sciences, 99*, 8378–8383.

Kunda, Z. (1999). *Social cognition*. Cambridge: MIT Press.

Kuo, B.-C., Stokes, M., & Nobre, A. (2012). Attention modulates maintenance of representations in visual short-term memory. *Journal of Cognitive Neuroscience, 24*, 51–60.

Lepsien, J., Thornton, I., & Nobre, A. (2011). Modulation of working-memory maintenance by directed attention. *Neuropsychologia, 49*, 1569–1577.

Levine, D., Warach, J., & Farah, M. (1985). Two visual systems in mental imagery: dissociation of "what" and "where" in imagery disorders due to bilateral posterior cerebral lesions. *Neurology, 35*, 1010–1018.

Li, W., Moallem, I., Paller, K., & Gottfried, J. (2007). Subliminal smells can guide social preferences. *Psychological Science, 18*, 1044–1049.

Lovett, M., Reder, L., & Lebiere, C. (1999). Modeling working memory in a unified architecture: An ACT-R perspective. In A. Miyake & P. Shah (Eds.), *Models of working memory*. New York: Cambridge University Press.

Mahon, B., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology Paris, 102*, 59–70.

McDowell, J. (1994). *Mind and world*. Cambridge: Harvard University Press.

Mikels, J., Reuter-Lorenz, P., Beyer, J., & Fredrickson, B. (2008). Emotion and working memory: Evidence for domain-specific processes for affective maintenance. *Emotion, 8*, 256–266.

Mitchell, K., & Johnson, M. (2000). Source monitoring: Attributing mental experiences. In E. Tulving & F. Craik (Eds.), *The Oxford handbook of memory*. Oxford: Oxford University Press.

Müller, N., & Knight, R. (2006). The functional neuroanatomy of working memory: Contributions of human brain lesion studies. *Neuroscience, 139*, 51–58.

Newell, A. (1990). *Unified theories of cognition*. Cambridge: Harvard University Press.

Oliveri, M., Turriziani, P., Carlesimo, G., Koch, G., Tomaiuolo, F., Panella, M., et al. (2001). Parietal-frontal interactions in visual-object and visual-spatial working memory: Evidence from transcranial magnetic stimulation. *Cerebral Cortex, 11*, 606–618.

Pasternak, T. & Greenlee, M. (2005). Working memory in primate sensory systems. *Nature Reviews Neuroscience, 6*, 97–107.

Paulescu, E., Frith, C., & Frackowiak, R. (1993). The neural correlates of the verbal component of working memory. *Nature, 362*, 342–345.

Peacocke, C. (1986). *Thoughts*. Oxford: Blackwell Press.

Peacocke, C. (1992). *A study of concepts*. Cambridge: MIT Press.

Postle, B. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience, 139*, 23–38.

Postle, B., Ferrarelli, F., Hamidi, M., Feredoes, E., Massimini, M., Peterson, M., et al. (2006). Repetitive transcranial magnetic stimulation dissociates working memory manipulation from retention functions in the prefrontal, but not posterior parietal, cortex. *Journal of Cognitive Neuroscience, 18*, 1712–1722.

Prinz, J. (2012). *The conscious brain: How attention engenders experience*. Oxford: Oxford University Press.

Santos, L., Barnes, J., & Mahajan, N. (2005). Expectations about numerical events in four lemur species. *Animal Cognition, 8*, 253–262.

Schacter, D. (2001). *The seven sins of memory*. New York: Houghton Mifflin.

Schacter, D., Addis, D., & Buckner, R. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience, 8*, 657–661.

Schacter, D., Addis, D., & Buckner, R. (2008). Episodic simulation of future events: Concepts, data, and applications. *Annals of the New York Academy of Sciences, 1124*, 39–60.

Schnall, S., Haidt, J., Clore, G., & Jordon, A. (2008). Disgust as embodied moral judgment. *Personality and Social Psychology Bulletin, 34*, 1096–1109.

Schwarz, N., & Clore, G. (1983). Mood, misattribution, and judgments of well-being: Informative affective states. *Journal of Personality and Social Psychology, 45*, 513–523.

Schwarz, N., & Clore, G. (2003). Mood as information: 20 years later. *Psychological Inquiry, 14*, 296–303.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience, 8*, 1391–1400.

Shanahan, M. (2010). *Embodiment and the inner life*. New York: Oxford University Press.

Shergill, S., Brammer, M., Fukuda, R., Bullmore, E., Amaro, E., Murray, R., et al. (2002). Modulation of activity in temporal cortex during generation of inner speech. *Human Brain Mapping, 16*, 219–227.

Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Oxford: Blackwell.

Sreenivasan, K., Sambhara, D., & Jha, A. (2011). Working memory templates are maintained as feature-specific perceptual codes. *Journal of Neurophysiology, 106*, 115–121.

Stanton, G., Bruce, C., & Goldberg, M. (1995). Topography of projections to posterior cortical areas from the macaque frontal eye fields. *Journal of Comparative Neurology, 353*, 291–305.

Todd, J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature, 428*, 751–754.

Tye, M. (1995). *Ten problems of consciousness*. Cambridge: MIT Press.

Uddin, L., Supekar, K., Amin, H., Rykhlevskaia, E., Nguyen, D., Greicius, M., et al. (2010). Dissociable connectivity within human angular gyrus and intraparietal sulcus: Evidence from functional and structural connectivity. *Cerebral Cortex, 20*, 2636–2646.

Vuontela, V., Rama, P., Raninen, A., Aronen, H., & Carlson, S. (1999). Selective interference reveals dissociation between memory for location and color. *NeuroReport, 10*, 2235–2240.

Winkielman, P., Berridge, K., & Wilbarger, J. (2005). Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin, 31*, 121–135.

Young, R., & Lewis, R. (1999). The Soar cognitive architecture and human working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory*. New York: Cambridge University Press.

Zwaan, R., Stanfield, R., & Yaxley, R. (2002). Do language comprehenders routinely represent the shapes of objects? *Psychological Science, 13*, 168–171.