Word Counts: Abstract: 246; Main text: 13888; References: 2007; Entire text: 16318

From format to function: Embodiment and the functional roles of neural circuits

Tyler Marghetis [a]
(tmarghet@ucsd.edu)

Tristan Davenport [a]
(trdavenp@ucsd.edu)

Ross Metusalem [a]
(rmetusal@ucsd.edu)

Benjamin K. Bergen [a]
(bkbergen@ucsd.edu)

[a] Department of Cognitive Science, UC San Diego, La Jolla, CA, USA, 92093-0515

Corresponding author:
Tyler Marghetis
Department of Cognitive Science
University of California, San Diego
La Jolla, CA 92093-0515
Email: tmarghet@ucsd.edu
Phone: 619-252-7798
Fax: 858-534-1128

Short Abstract

When it comes to specifying the representational format of cognition, so-called "embodied" or "modal" approaches have gained prominence. Empirical evidence suggests that systems for perception and action are involved in higher cognition, but critics have challenged the embodied approach and defended the role of amodal representations. We argue that questions of representational format are unresolved by current evidence—and that, moreover, the debate is irresolvable *in principle*. For both empirical and theoretical reasons, the widely assumed distinction between modal and amodal representations is untenable. Instead, we defend an approach that focuses on the functional role of functionally specialized neural circuits.

Abstract

There remains substantial disagreement about the representational format of higher cognition. In recent years, so-called "embodied" or "modal" approaches have gained prominence, arguing that higher cognition is grounded in brain areas dedicated to perception and action. Indeed, empirical evidence suggests that higher cognition might rely on brain areas classically viewed as perception and action systems. This has led some to argue that the representational format of higher cognition is embodied or modal. However, critics have questioned the sufficiency of an embodied approach and have defended the role of amodal processes and representations, perhaps working in concert with modal representations grounded in perception and action. In this paper, we argue that questions of representational *format* are left unresolved by current evidence (Section 3). Moreover, we argue that the debate cannot even be resolved *in principle*. For both empirical and theoretical reasons, the widely assumed distinction between modal and amodal processing or representation is untenable at the neural level (Section 4). We propose that more productive than trying to uncover the representational format of cognition would be to address the *functional roles* of functionally specialized neural circuits (Section 5). This approach retains an interest in the functional role of brain areas known to support perception and action, without reifying these areas as different in kind from "amodal" brain areas. We end by identifying desiderata for experimental design, pointing out promising avenues of research, and discussing the benefits of focusing on function rather than format (Section 6 and 7).

Keywords: embodiment; simulation; perception; action; representational format; functional role

1. Introduction

What is the representational format of higher cognition? Einstein famously claimed to reason almost entirely in visual or motor images (Hadamard, 1954), but concepts like "justice" or "the number two" seem to demand rather different representational resources, more abstract than sensorimotor. Central to cognitive science since its inception, questions of representational format follow naturally from an interest not merely in the physical mechanisms that instantiate cognition—the implementation—but also in functional- and computational-level accounts of cognition (Marr, 1982). The reasoning goes that this functional machinery of cognition has to be realized in some particular format. Carey (2009) gives voice to this stance when she writes, "A full characterization of any mental representation *must specify its format* as well as its conceptual role" (p. 457; emphasis added). Indeed, much of contemporary cognitive science has been devoted to investigating the nature of mental representation (e.g. Minsky, 1977; Kosslyn, 1980; Johnson-Laird, 1983; Barsalou, 1999; Glenberg & Robertson, 2002; Pylyshyn, 2002; Paivio, 2007).

Yet there remains substantial disagreement about the representational format of higher cognition. One particularly fierce debate has asked whether representations are "embodied" and "modal," or "abstract" and "amodal." Proponents of the amodal approach argue that the representations and processes of higher cognition differ in format from the perceptual modalities that activate them and the motor outputs that they can drive (Fodor, 1975; Kintsch, 1998, Newell & Simon, 1972). Proponents of the modal view, in contrast, suggest that at least some representations and processes of higher cognition are not entirely abstracted from perception or action, but have an inherent perceptuo-motor nature. This view is not entirely novel (Hume, 1739; Locke, 1690; see Prinz, 2002), but has experienced a recent resurgence among cognitive scientists (Lakoff, 1987; Johnson, 1987; Barsalou, 1999; Glenberg & Robertson, 2000; Paivio, 2007; Bergen, 2012; among others).

The debate over representational format has generated a wide range of research on reasoning, categorization, language comprehension and production, and mathematical cognition, among other topics (Barsalou, 2010). Some pro-embodiment authors characterize their experimental results as specifically conflicting with amodal views of representation (e.g. Glenberg & Robertson, 2000); others make the more modest point that their results are predicted by modal but not amodal

theories (e.g. Stanfield & Zwaan, 2001). The latter rhetoric is more common. Consider an early, typical example:

> "Although the current experiment cannot be taken as an unconditional falsification of amodal symbol systems [...], it does suggest that amodal theories, as they stand now, are insufficient to fully explain comprehension and that, moreover, perceptual symbol systems or, at the very least, hybrid systems, are a viable alternative." (Stanfield & Zwaan 2001, p.156)

Quite recently, there has a been some push-back in the literature, with several papers arguing that results purporting to support modal accounts really cannot distinguish between modal and amodal theories, since they are compatible with appropriately extended versions of either (e.g. Mahon & Caramazza, 2008; Dove, 2009; Chatterjee, 2010; Binder & Desai, 2011). These authors have proposed that modal theories of cognition are insufficient and need to incorporate both modal and amodal representations (e.g. Mahon & Caramazza, 2008; Dove, 2009) or admit gradations from modal to amodal symbols (Chatterjee, 2010; Binder & Desai, 2011). If existing empirical work fails to adjudicate between modal and amodal accounts, then these critiques have identified a significant failing.

In this paper, we argue not only that these criticisms are correct, but moreover that they do not go far enough. It's not merely the case that existing evidence cannot distinguish between modal and amodal representational formats. Rather, the representational debate rests on an unjustified premise—a distinction in kind between modal and amodal neural systems—and we argue that no type of neural evidence could *ever* resolve it. We propose instead that the field would be better served by focusing its attention on questions of mechanism and function. As it stands, there is substantial current evidence demonstrating that brain areas known to subserve perception and action are also activated during higher cognitive processes, such as language comprehension (see Section 3, below). We defend a research program that focuses on the functional role of these specialized brain circuits in higher cognition. That is, instead of asking about representational format, we argue that the question that ought to be driving the field is: Which neural circuitry, performing which functions for perception and action, also serves what functional role(s) in what higher cognitive processes (if any)?

The structure of this paper is as follows. First, we describe the modal and amodal positions in the representational format debate in more detail (Section 2). We then canvass the empirical evidence adduced by both sides of the debate, focusing on language research, a particularly contentious area (Section 3). We next argue that the debate is in principle irresolvable (Section 4), and suggest that the field would be well served by re-orienting towards the functional roles of functionally specialized neural circuits (Section 5). Finally, we outline some of the promising empirical avenues that may allow us to address this issue (Section 6), and conclude by discussing the potential rewards of this modified approach (Section 7). We focus throughout on language research, although the proposal applies broadly to questions of representational format and the neural basis of cognition.

2. The modal-amodal debate

A number of theorists have defended the idea that higher cognition is grounded in perception and action. Some of the most fully articulated accounts can be found in the work of Larry Barsalou (Barsalou, 1999, 2008, 2010; Barsalou et al., 2003; i.a.), Jesse Prinz (Prinz, 2002; i.a.), and George Lakoff and Mark Johnson (Johnson, 1987; Lakoff, 1987; Lakoff & Johnson, 1999, i.a.), so we draw heavily from their work here. These authors have been responding to the assumption, often implicit, that higher cognition relies on representations and processes that are amodal—that is, divorced from perception and action. The debate hinges, therefore, on questions of representational format: what it means for a representation to be modal, and whether the representations involved in higher cognition are of this sort. In this section, we briefly review the claims of both sides of the debate.

2.1 The modal position

Modal approaches argue that the representations and processes of higher cognition are based in modal systems—those dedicated to perceiving and acting. This is nicely illustrated by Barsalou's (1999) characterization of the relation between color perception and color concepts. Whether engaged in perception or cognition, Barsalou suggests, the representation of color involves the same neural substrate—in particular, the "neural systems that represent color in perception" (1999, p.578). More generally, the representations underlying higher cognition are grounded in perception because "they are represented in the *same systems as the perceptual states* that produced them" (p.578; our emphasis). Similarly, Prinz (2002), argues that "concepts derive from perceptual representations" (p.

113), and "are couched in representational codes that are *specific to our perceptual systems*" (p. 119; our emphasis). And Lakoff and Johnson (1999) write that an "embodied concept is a neural structure that is actually *part of, or makes use of, the sensorimotor system* of our brains." (p. 20; our emphasis). In short, all argue that at least some conceptual representations are modal in virtue of their reliance on perceptual systems in the brain.

Most modal theories are therefore explicitly reductionist, in the sense that they make claims about the biological substrate of cognition. It is a claimed overlap at the neural level that gives concepts their perceptual or motor format. Importantly, modal accounts often do not specify whether it is conceptual *representations* or *processes* that are proposed to be modal (Machery, 2007). See, for instance, the quote from Lakoff and Johnson, above, in which they equivocate between being "part of" and "mak[ing] use of" a sensorimotor system—an equivocation between representation and process. Arguably, this is precisely because of the reductionist move these theorists take; at the level of neurons, there are no precise homologues for the algorithmic-level distinction between representation and process.

Implicit in the pro-embodiment line of argument is a commitment to certain neural systems being devoted to perception and action—that is, "modal"—such that using these systems for other cognitive functions, such as language comprehension, makes those latter functions themselves modal. For instance, Barsalou and colleagues (2003) describe the neural grounding of concepts in perceptual systems by identifying the neural structures hypothesized to serve this function:

> Modal approaches represent knowledge very differently [from amodal approaches]. Rather than being transduced into amodal symbols, modality-specific states are captured by *adjacent memory systems*. Consider visual states (analogous accounts exist for states on other modalities). When a car is perceived visually, a set of neural feature detectors becomes active in the *visual system*. Conjunctive neurons in a nearby association area then conjoin the active features and store them in memory. Later, in the absence of visual input, these conjunctive neurons partially reactivate the original set of feature detectors to represent the car visually. Such re-enactments or simulations […] provide the cognitive-

> level representations that support memory, language and thought. (p. 85; our emphasis)

The basic claim of modal approaches, therefore, is that there are perceptual or motor structures in the brain that are also used for conceptual functions, and that this neural repurposing is what makes those conceptual functions *modal*. In sloganized form, it states:

> **The Modal Hypothesis**: Neural machinery that is specifically modal—that is, devoted to perception or action—also plays a functional role in higher cognitive processes.

The Modal Hypothesis is distinct from a number of superficially similar proposals. It is orthogonal, for instance, to the Piagetian claim that conceptual representations have their developmental origins in sensorimotor schemas. And it is stronger than the claim that mature concepts retain vestigial connections to sensorimotor systems, without those systems playing a functional role in conceptual processing. Moreover, the contemporary modal-amodal debate is *not* the same as the "Imagery Debate," which asked whether imagery relies on analog or propositional representations (e.g. Kosslyn, 1980, 1994; Pylyshyn, 2003). Adherents to the Modal Hypothesis can, and often do, remain agnostic as to the implications for the classic imagery debates (Nichols et al, 1996; Thomas, forthcoming). For example, while Barsalou (1999) explicitly stated that modal representations are analogical, more recent writings have dropped this claim and instead focus on the contribution of perceptual brain systems (Barsalou, 2008, 2010). Indeed, Pylyshyn (2002) argued that one can accept that perceptual brain systems are involved in imagery, while maintaining that all representations are symbolic. Contemporary versions of the Modal Hypothesis, therefore, are specifically about the role of sensorimotor brain systems in higher cognition.

The Modal Hypothesis, however, has received considerable opposition. We turn now to those criticisms.

2.2 The amodal position

Modal accounts often provide clear characterizations of what they take the amodal position to be and its differences from the modal position. For instance, Barsalou and colleagues (2003) contrast

the amodal approach with the modal one in terms of the neural systems responsible for cognitive processes and representations:

> In amodal approaches, sensorimotor representations are transduced into an amodal representation [….] that *reside[s] outside sensorimotor systems* […]. Once amodal redescriptions of sensorimotor states exist, all cognitive processes operate on them to achieve their functions — not on memories of the original sensorimotor states. (p. 85; our emphasis)

So when Barsalou (1999) describes amodal representations as "inherently nonperceptual" (p. 578) he means in part that amodal symbols and perceptual representations do not use the same neural systems:

> The amodal symbols that represent the colors of objects in their absence reside in a *different neural system* from the representations of these colors during perception itself. (578, our emphasis)

While characterizations of the amodal thesis are commonplace in the writing of its opponents, it's harder to find explicit advocacy of amodal conceptual representations, perhaps because the view is widely assumed. But such characterizations do exist. Critically evaluating the Modal Hypothesis, Mahon and Caramazza (2008, p. 60) state that the amodal position holds that concepts are "an 'abstract' and 'symbolic' representation that would not be constituted by information in the sensory and motor systems."

The amodal symbols hypothesis, therefore, holds that non-perceptual processing is not performed by perceptual systems, but by amodal processes and representations that live outside the perceptual and motor systems. At first pass, this amodal hypothesis seems to contrast radically with the modal position. We turn now to the evidence that is thought to adjudicate between these two views.

3. Empirical evidence for and against modal accounts of language use

A large literature has emerged over the past decade in support of the Modal Hypothesis, purporting to demonstrate that higher cognition is based in perception and action (see reviews in Barsalou 2008, 2010). This includes work on recall, imagery, reasoning, and arithmetic, among others, but likely the richest segment of this literature addresses language comprehension. Studies have demonstrated that

processing language about action results in engagement of the motor system (e.g. Glenberg, et al., 2008), and that processing language about perceptual experience engages perceptual systems (e.g., Kaschak et al, 2005). This work is often taken to show that the semantic representations built up during language comprehension are modality-specific and tied to perceptual and motor systems. In this section, we briefly review a sample of prominent studies in this area. We then discuss evidence against—and general skepticism about—the view that language comprehension relies on modal semantic representations.

3.1 Engagement of modal neural systems during language processing

A broad range of experimental techniques in cognitive neuroscience and psychology have been used to investigate the engagement of modal neural systems during language comprehension, including reaction time paradigms, functional magnetic resonance imaging (fMRI), and transcranial magnetic stimulation (TMS). Generally speaking, each technique allows the researcher to infer whether comprehension of language about perception or action interacts with actual action or perception. When such results are found, researchers often conclude that perceptual and motor brain systems play a functional role in language comprehension, and that language comprehension is therefore modal.

A great deal of work in this area has used behavioral paradigms. Reaction time studies typically pair a linguistic stimulus with a nonlinguistic task—either perceptual or motor—and require a speeded behavioral response to either the linguistic or nonlinguistic stimulus. If the compatibility of the linguistic and non-linguistic stimuli has a significant effect on reaction time, this is taken as evidence that language processing has engaged modal brain systems. A seminal study by Glenberg and Kaschak (2002), for instance, showed that comprehending sentences implying action toward (e.g., *Courtney handed you the notebook)* or away from the body (e.g., *You handed Courtney the notebook*) selectively affected the time required to launch subsequent arm movements toward or away from the body. In the study, participants held down a central button to trigger visual presentation of a sentence and were asked to make a yes/no sensibility judgment. To respond, participants released the central button and launched an arm movement toward or away from the body in order to press either the "yes" or "no" button, with response direction counterbalanced within subjects. As predicted, after comprehending an "away" sentence, participants were faster to launch an arm movement away from the body than toward the body; for "toward" sentences, this pattern reversed.

This "action-sentence compatibility effect" held for both concrete (e.g., *You handed Courtney the notebook.)* and abstract sentences (e.g., *You told Liz the story.*). The authors concluded that language comprehension requires perceptual-motor simulation in neural systems responsible for perception and action, and "real bodily action is at the root of meaning conveyed by language" (p. 563). Subsequent behavioral studies have further explored the nuanced and dynamic ways in which the motor system is engaged during language comprehension (e.g., Bergen & Wheeler, 2010; Masson et al., 2008; Zwaan & Taylor, 2006).

Reaction time studies have also suggested the engagement of perceptual systems during language comprehension, with prominent examples coming from studies by Rolf Zwaan and his colleagues. In this work, participants hear sentences implying specific perceptual features of some critical object. After each sentence, subjects see a picture of an object and judged whether this object was mentioned in the sentence. In a series of studies that have looked at implied object orientation (Stanfield & Zwaan, 2001), shape (Zwaan et al., 2002), and visibility (Yaxley & Zwaan, 2007), Zwaan and colleagues have observed faster responses when the relevant perceptual features of the picture matched those implied by the sentence (e.g., faster responses to a picture of an eagle with outstretched rather than folded wings following the sentence *The ranger saw the eagle in the sky*). Such results are taken to support the notion that "the representation of meaning from linguistic input is a dynamic process involving malleable perceptual representations rather than the mechanical combination of discrete components of meaning" (Zwaan et al., 2002, p. 170). Other behavioral studies have addressed the perceptual basis of linguistic representations of color (Connell, 2007), literal motion (Kaschak et al, 2005) and fictive motion (Matlock, 2004), and the spatial representations of concrete and abstract verbs such as *respect* and *push* (Richardson et al., 2003), to name a few.

Further support has come from neuroimaging studies that find activation in perceptual-motor systems during language comprehension. For instance, comprehending language about effector-specific actions (e.g. actions using the hand, foot, or mouth) selectively activates areas of motor and premotor cortex responsible for observing and executing actions with those effectors (Aziz-Zadeh et al., 2006; Hauk et al., 2004; Tettamanti et al, 2005). In one study, Tettamanti et al. (2005) presented participants with sentences conveying actions of the hand, mouth, or leg, and found that these sentences, compared to abstract control sentences, activated a left hemisphere fronto-parieto-

temporal circuit known to be activated during action observation and execution. While all three effectors exhibited some overlap in activation, sentences describing effector-specific actions selectively activated regions responsible for the motor representation of that effector. The authors concluded that "understanding sentences conveying an action-related content requires the contribution of sensorimotor circuits, partially overlapping with those active during the execution and observation of the same actions" (p. 278). Similar results have been observed, and similar claims made, regarding language about visual and figurative motion, and brain areas dedicated to motion perception (MT+) (Saygin et al., 2010).

In addition to reaction time and fMRI studies, recent studies using transcranial magnetic stimulation (TMS) have targeted the modality-specificity of language comprehension. TMS is a technique in which an electromagnetic field is generated at the skull, altering the functioning of a focused cortical region with a spatial resolution as precise as a few millimeters (Bolognini & Ro, 2010). TMS pulses to primary motor cortex, for instance, cause motor evoked potentials (MEPs) in the muscles of the associated effectors. Crucially, the amplitude of these MEPs is increased when there is a match between the site of stimulation and the semantic content of language stimuli (Buccino et al., 2005; Oliveri et al., 2004). More controversially, TMS has also been shown to modulate overt responses to linguistic stimuli, which some argue is stronger evidence for a functional role of modal neural circuits in language comprehension, since the effect is on language comprehension itself (Willems et al., 2011). Pulvermüller and colleagues (2005), for example, had participants perform a lexical decision task in which critical words referred to either leg actions (e.g., *kick*) or hand/arm actions (e.g., *pick*). A single subthreshold TMS pulse was applied to hand and leg areas of the left hemisphere 150 milliseconds after the onset of the critical words. Pulvermüller and colleagues found that TMS to the arm area resulted in faster responses to arm than leg words, and conversely that TMS to the leg area resulted in faster responses to leg than arm words. Willems and colleagues (2011) report similar results using repetitive TMS (rTMS) on premotor cortex. These authors conclude that, since the stimulation of a motor area affects lexical decision times, the results establish a causal role for the motor system in language comprehension—which, they argue, is not necessarily demonstrated by behavioral and fMRI studies that find effects of language comprehension on subsequent action, or co-activation of language and motor areas.

Taken together, these and similar studies attempt to establish overlap between the neural systems responsible for perception and action and those underlying language comprehension. In many cases, the results are taken to indicate a perceptual-motor basis of the semantic representations constructed during language comprehension, in line with theories that place modal representations at the heart of higher cognition (Barsalou, 1999; Bergen & Chang, 2012). These researchers frequently posit that language comprehension involves embodied simulations of described situations, and that these simulations (and hence, language comprehension) rely on the motor and perceptual neural systems connected to the appropriate effectors and sense organs. The proposal, therefore, is that language comprehenders engage a distributed, modal semantic system, the components of which are housed in the same brain areas responsible for the various modalities of perception and motor control.

3.2 Response to the modal evidence

The evidence surveyed above points to a view that, at its most extreme, takes representations in perceptual-motor systems to exhaust the semantic representations built up during language processing. But there are good reasons to doubt the sufficiency of perceptual-motor systems. Mahon and Caramazza (2008), for instance, argue that modality-specific systems cannot account for all conceptual knowledge, pointing to highly abstract concepts like "incredulous" or "astute" as examples of concepts that lie outside the dominion of sensorimotor representation. Dove (2009: 414-415) articulates the general argument:

> All in all, the available evidence suggests that some of our concepts employ perceptual representations—particularly concrete or highly imageable concepts—but fails to support the conclusion that perceptual symbols are the lingua franca of concepts.

Beyond the objection that there does not exist direct evidence for the sufficiency of modal representations, this extreme version of the Modal Hypothesis is also challenged by evidence that conceptual knowledge can be impaired in the absence of associated sensorimotor impairments. For instance, in *conceptual* or *ideational* apraxia, caused by lesions to the parietal-temporal-occipital junction, patients fail to use tools in contextually appropriate ways (e.g. brushing teeth with a comb), but nevertheless remain able to make the fluid hand motions required for tool use (Ochipa et al., 1989).

But it's important to note that these objections only attack the sufficiency of modal representations, a hypothesis that does not appear in the pro-modal literature. It is true that there are proposals on the table for grounding even the most abstract concepts in perception and action (Prinz, 2002, ch. 7; Lakoff & Johnson, 1980; Gallese & Lakoff, 2005; Pezzulo et al., 2011). But crucially, the Modal Hypothesis does not deny that language or other higher cognitive functions are subserved by multifarious networks of brain regions, engaged dynamically and opportunistically. Any serious account of access to word meanings or concepts will include roles of varying importance for recall, planning, social cognition, affect, and a host of other capacities. For this reason, criticisms of modal accounts that target sufficiency are akin to noting that lesions to V4 affect vision, and concluding that primary visual cortex is not sufficient for object perception. Naturally, primary visual cortex isn't sufficient for vision, and the interesting vision science asks not whether any parts of the vision system are sufficient, but rather *what the different parts do*. Even the staunchest defenses of modal representations (Barsalou et al, 2008; Prinz, 2002, ch. 7) support the involvement of mechanisms that are not properly modal, including symbolic linguistic encoding. Attacks that target sufficiency thus seem misguided, since the spirit of the Modal Hypothesis focuses not on sufficiency but on the widespread re-deployment of modal brain areas.

Potentially more troublesome for the Modal Hypothesis are arguments that target the *necessity* of modal processes and representations. If modal systems are truly necessary for conceptualization or language comprehension, then semantic knowledge about perceptual and motor concepts should be impaired whenever perceptual and/or motor systems are impaired—and thus, we should not find a preservation of semantic knowledge in cases where corresponding perceptual or motor knowledge is impaired. And yet, dissociations like this have been observed, for instance, dissociations between semantic and motor representations of tool use (for review, see Johnson-Frey, 2004). Patients with *ideomotor* apraxia—caused by lesions to the left posterior parietal cortex or the left premotor cortex—retain context-sensitive knowledge of tools' functions but are unable to perform the appropriate actions with those tools (Geschwind & Kaplan, 1962; Sirigu et al., 1995). Such dissociations have been taken as evidence that modality-specific representations of relevant motor skills are not necessary for conceptual knowledge of action (Mahon & Caramazza, 2008).

Unlike sufficiency, the necessity of perceptual-motor involvement in higher cognition is a genuine concern for modal accounts that argue for a functional role for modal systems. Note, however, that

the Modal Hypothesis does not, on its own, entail that perceptual-motor systems are *always* necessary, only that—when available and appropriately engaged—those systems make specific contributions to, and thus are partially constitutive of, language processing and other facets of higher cognition. In Section 5 below, we argue that a productive account of language and higher cognition should focus on such a mechanistic approach, and thus abandon concerns with the blanket necessity or sufficiency of sensorimotor representations.

There are other serious objections to the Modal Hypothesis. Mahon and Caramazza (2008) have pointed out that many, if not all, of the experimental results supporting a modal conceptual system are compatible with a completely amodal conceptual system equipped with associative links to perceptual and motor representations. In sloganized form, their argument is as follows:

> **The Mediating Process Reply:** The observed relation between sensorimotor processing and higher cognitive functioning is due not to shared representations or processes, but to downstream activation spreading between amodal and modal areas.

According to this Mediating Process Reply, the activation of modal systems by word stimuli, for example, is the result of first activating the associated amodal concept in an amodal system, followed by activation spreading via associative links that mediate between amodal and sensorimotor processes. On this account, sensorimotor systems themselves do not play a causal role in higher cognition; their engagement during conceptual processing is epiphenomenal.

Versions of this objection have been voiced elsewhere (Chatterjee, 2010; Dove, 2009). It is important to note that this argument applies as much to reaction time results as it does to fMRI (Pulvermüller et al., 2005) and some TMS results (Willems et al., 2011). As Mahon and Caramazza (2008) point out, for instance, amodal accounts can even accommodate the facilitation of lexical decisions following TMS to sensorimotor areas (Pulvermüller et al., 2005, Willems et al., 2011), since it is entirely possible that TMS-induced motor activity can spread to purely amodal concepts, activating that amodal representation and producing the observed facilitation of lexical decision times. None of the evidence described in Section 3.1, therefore, can be straightforwardly interpreted as showing a causal or functional role for modal neural systems in conceptual processing.

In broad strokes, then, that is the current state of the field. The modal-amodal representation debate is at something of an impasse. It is clear that sensorimotor tasks can affect semantic processing, and that semantic processing can engage sensorimotor systems. However, it is unclear what implications these results may have for the format of mental representations, and it also remains unclear what functional role this activity might have in higher cognition, if any. To further complicate matters, there is little agreement on what kind of evidence could in principle resolve this issue. Many theorists look to neuroimaging for the crucial evidence, yet it seems that any neuroimaging finding that purports to support the Modal Hypothesis could also be explained by a post-hoc amodal Mediating Process Reply. In the next section, we suggest that this impasse is not merely a shortcoming of ingenuity or experimental technology, but a sign that the modal-amodal distinction is untenable at the neural level.

4. The modal-amodal distinction is untenable

Debates surrounding the Modality Hypothesis have largely assumed that the opposing positions are meaningful, coherent, and distinct—that the debate is resolvable in principle, if as yet unresolved. In this section, we challenge this assumption. We consider the criteria used to determine when a brain system is "modal," and argue that at the neural level this pursuit of labels is quixotic. If questions of representational format are to be adjudicated at the neural level, as they are in the literature concerned with the Modal Hypothesis, then a necessary first step is articulating criteria that can reliably pick out a "modal" neural system. Existing definitions, however, preclude the possibility of such a system playing a role in higher cognition, or fail to characterize the actual activity of those neural systems we intuitively consider "perceptual" or "motoric." In short, when framed at the neural level, debates about the involvement of modal systems lack an explicit definition of what constitutes a modal system. In light of this conclusion, we argue that while recent proposals to marry modal and amodal views do take steps in the right direction, ultimately we must move beyond the notion of modality-specificity when studying how the brain accomplishes higher cognitive tasks.

4.1 The brain flouts the modal-amodal distinction

As discussed in the previous section, the current representational debate appeals explicitly to the brain in determining whether cognitive representations are modal or amodal. Argumentation generally takes the following form: "Engaging representation X activates modal brain system Y;

therefore, representation X must be modal." The argument pivots entirely on the word "modal." *But without a definition that demarcates between modal and amodal neural systems, one cannot appeal to the brain to determine the modality-specificity of cognitive processes or representations.* As Aydede (1999) argued in a commentary on Barsalou's (1999) original BBS target article, for us to make claims about the perceptual grounding of cognition, "we must have a workable notion of perceptual/nonperceptual systems," (p. 610)—and, Aydede notes, Barsalou (1999) does not supply one beyond the definition that concepts are modal when they "are represented in the same systems as the perceptual states that produced them" (578).

In a response to Aydede's criticism, Barsalou (1999) appealed to common sense notions of perceptual systems. We quote at length because this gets to the heart of our argument:

> Quite remarkably, Aydede claims that there is no way to distinguish perceptual and nonperceptual systems independently of modal and amodal symbols. Given the huge behavioral and neural literatures on perception, this is a rather surprising claim. A tremendous amount is known about the perceptual systems of the brain independent of anything anyone could possibly say about conceptual systems. Clearly, we know a lot about perceptual abilities and the neural mechanisms that implement them. *Identifying these abilities and mechanisms independently of perceptual symbols is not only possible but has long since been accomplished in the scientific disciplines that study perception. In defining perceptual symbols, I simply used these well-recognized findings.* [….] Perceptual symbols *follow the well-trodden paths of perception researchers* in claiming that the representational mechanisms of perception are the representational mechanisms of knowledge. *What is perceptual about perceptual symbol systems is perception in the classic and well-established sense.* (1999, p.640, emphasis ours)

Whether this "classic and well-established sense" is sufficiently coherent to sustain the "grounding" of conception in perception, however, is the very issue that Aydede raises. Without a precisely articulated account of when a neural circuit qualifies as "modal"*,* one can only appeal to convention and common sense—notoriously unreliable arbiters of neural architecture.

Since then, few have responded to Aydede's challenge. Prinz (2002) offers one of the clearest and most cogent responses, a careful account of when exactly a particular neural population counts as a perceptual system or *sense*[1]. He characterizes perceptual systems as *dedicated input systems,* which he defines as follows:

> In saying that senses are *input* systems, I mean that they are systems that **receive inputs from outside of the brain**. […] (116, bold emphasis ours)
>
> In saying that the senses are *systems*, I want to emphasize the fact that they each consist of their own sets of operations and representations, housed in **separate neural pathways. Distinguishing separate neural pathways is crucial**. […] To say that senses are **systems** means that they can be **divided up internally, in our case, by distinct collections of cooperative neural populations**. (115-116, bold emphasis ours)
>
> Finally, in saying that senses are *dedicated*, I mean that each sense responds to a **proprietary input class**." (117, bold emphasis ours)

So a dedicated input system takes a particular class of outside input (e.g. visual input), and processes it in a specific and exclusive neural pathway (e.g. the visual pathway, starting at, say, V1). This position is certainly appealing for its definitional rigor and simplicity, and captures the received wisdom that particular brain regions are responsible for each perceptual modality, others are responsible for action, and so on. Yet, in the light of the very results on conceptual processing that they are recruited to explain, Prinz's criteria appear to be untenable. Let us consider each in turn.

For a neural system to count as perceptual, according to Prinz's first criterion, it should receive inputs from outside the brain. But there is no principled way to terminate the causal chain that begins with sensory input and spreads rapidly throughout the brain. Indeed, the whole brain satisfies this first criterion, including systems that nobody has ever characterized as "perceptual"—systems like Broca's area and all the rest of the canonical language machinery. Second, the candidate neural system should consist of "separate neural pathways," "distinct collections of cooperative neural populations." But again, this does little to distinguish putative perceptual systems from other specialized neural populations, also made up of neurons that "work in tandem" but not in the

service of perception. So the first two criteria do little to parcel off the perceptual systems from the rest of our neural machinery.

That brings us to the third criterion, that a perceptual system needs to be dedicated—that is, to respond to a proprietary input class. And here's the crux of our argument: We know that neurons, even in primary perceptual areas, respond not merely to afferent input from specific perceptual organs (e.g. the eyes), but also to efferent input from other cortical sources. After all, this is precisely the finding that drives modal theories in the first place: that classical perceptual and motor areas are also engaged during higher cognitive functions like recall, imagery, language processing, and so on. What's more, the activity of classical perceptual areas "dedicated" to one perceptual sense can be modulated by input to other sensory modalities (Ghazanfar & Schroeder, 2006; Kayser & Logothetis, 2007). For instance, both primary auditory cortex and primary somatosensory cortex have been shown to respond to stimuli in the other modality: activity in A1 is modulated by tactile stimuli, and activity in S1 by auditory stimuli (Lemus et al., 2010). Findings like these led Ghazanfar and Schroeder (2006) to conclude that "...much, if not all, of neocortex is multisensory. This necessarily forces us to reconsider the validity of probing the brain unimodally and suggests a different perspective when considering other aspects of cognition [...]." (p. 284). We agree. Neurons in "perceptual systems" do more than just process perceptual input, let alone input from one particular "proprietary input class." With multisensory processing occurring even in primary sensory cortices, and putative "perceptual" systems responding to rampant cortico-cortical connections, it is impossible to define specific neural systems, at least cortical ones, as instantiating a specific perceptual modality using a criterion that appeals to proprietariness of input. After all, it is this very input promiscuity that modal accounts are trying to explain.

In essence, this difficulty in applying Prinz's third criterion boils down to the question of what it means for a neuron or neural population to respond to input from apparently dissimilar sources. The study of "mirror neurons" is illuminating in this regard. There, the reasoning is perhaps clearer. Suppose a neuron fires both when an animal performs a specific action and when it observes another animal performing the same action (Rizzolatti & Craighero, 2004). What type of cell is this, and how should we characterize its behavior? Is it a motor cell, being used for perception as well—that is, are perceptual representations somehow "motoric"? Is it a perceptual cell, used for motor control as well—that is, are motor representations somehow "perceptual"? Is it a multimodal cell,

used for perception *and* action, or an amodal cell, encoding the action in the abstract? The mirror neuron literature avoids such debates about the format of mirror neuron representation—mere notational negotiations—and focuses instead on the functional contributions of the mirror neuron network (e.g., Rizzolatti & Sinigaglia, 2010).

This same logic applies to "perceptual" and "motor" systems that are also activated during non-perception or non-action. Consider for the sake of argument a hypothetical cell that responds to a particular range of wavelengths of light, but also becomes active when a person comprehends a description of color in that range. How should we characterize such a cell? Is it a perceptual cell also used for conceptualization? A conceptual cell also used for perception? A multifunctional cell used for perception and conception? Or an amodal cell that encodes color, independently of whether it's being perceived or imagined? At the neural level, these sorts of labels dissolve because they are not meaningfully distinct.

To summarize, we've argued that at the neural level, there is no principled distinction between "modal" brain systems and "amodal" ones. Responding to Prinz's three criteria, we showed that the (i) the *whole* brain "receives inputs from outside the brain," (Prinz, 2002, p. 116) (ii) the *whole* brain is characterized by "collections of cooperative neural populations," (ibid) and (iii) *no part* of the brain—not even primary sensory cortices—is truly response to only one "proprietary input class" (ibid, p. 117). Indeed, the last criterion—proprietariness of input—goes against the central insight of the Modal Hypothesis: that brain areas involved in perception *can be co-opted to do work besides perception*, which necessarily involves processing non-perceptual inputs.

This rampant multifunctionality of perceptual and motor neural systems undermines the very notion of a "modal" neural system. While we have focused on the account offered in Prinz (2002), we have done so only because he has so clearly articulated the assumptions inherent in any attempt at neural topography that hopes to locate modality-specific systems. Neural considerations cannot adjudicate between the possible labels we might want to apply to a mirror neuron—whether motor, perceptual, multimodal, or amodal—any more than to neurons involved in both perception/action and conception. And that's the rub. When we look closely at how it maps to the neural level, the modal-amodal debate—while intuitively compelling—reduces to a terminological quibble. Specifically,

when the inputs to a neural system are heterogeneous, as we have argued they are, the distinction between modal and amodal format dissolves at the level of neurons and their functional properties.

4.2 A word about function and teleology

Of course, there may be ways of determining representational format that do not appeal exclusively to neural implementation. Teleological accounts of mental content, for instance, may hold promise for distinguishing modal and amodal systems (Millikan, 1984; cf. Anderson 2010, p. 266). From the perspective of natural selection, the function of a particular organ—or neural region—is determined in light of that organ's history, its contribution to fitness over evolutionary time. As Sober (1993, p. xi) notes, "adaptation is a historical concept." Just as the heart's function is to pump blood, not to make a lub-dub noise, one might argue on evolutionary grounds that MT+ is modal because, over evolutionary time, its adaptive function was the perception of visual motion.

We are friendly to the spirit of this project. But it is an entirely different project from the one engaged in by theorists on both sides of the Modal Hypothesis. As underscored by our review of arguments for and against the Modal Hypothesis, the target of the current debate, in theory and in practice, is the use and re-use of a particular family of neural systems. The target is not the *historical origins* of the present-day computational properties of a particular neural region—whether over developmental time, for those of empiricist bent, or over evolutionary time for nativists. The concern is the contribution of particular cortical regions *today*, in anatomically modern human brains—more specifically, the possibility that a single cortical region could make contributions to multiple cognitive processes, perhaps as disparate as, say, motion perception and the comprehension of language about motion (e.g. Kaschak et al., 2005). Indeed, both proponents and opponents of the Modal Hypothesis seem to have a tacit agreement about which neural populations should count as modal, notwithstanding the lack of coherent criteria for picking out such populations. Appeals to teleology or adaptive function, therefore, are not germane.

4.3 "Hybrid" positions

In section 3.2, we argued that much of the current evidence does not adjudicate between modal and (sufficiently elaborated) amodal accounts. We are certainly not the first to recognize this, and other authors have responded by proposing "hybrid" accounts that wed modal and amodal representations (Mahon & Caramazza, 2008; Davies, 2004; Markman & Dietrich, 2000; Dove, 2009;

Chatterjee, 2010; Binder & Desai, 2011; Meteyard et al., 2012). Markman and Dietrich (2000), for instance, argue that the amodal view should be supplemented rather than discarded, so that "cognitive models [become] more sensitive to perceptual representation" (p. 475). An illustrative example of a hybrid account is Mahon and Caramazza's (2008) *Grounding by Interaction* hypothesis. On their account, a concept has stable and context-independent amodal content, but is also enriched by perceptual-motor content that supports its use in a particular context, at a particular point in time, in a contextually appropriate way. Similarly, Dove (2009) defends a pluralistic view in which separate modal and amodal semantic codes subserve processing of concepts of high and low imageability, instantiated by distinct neural systems. Most recently, authors like Binder and Desai (2011) and Chatterjee (2010) have proposed a *spectrum* of modality. For instance, Chatterjee (2010) argues that the brain has neuroanatomical axes along which representations vary from unimodal to amodal:

> The nervous system has unimodal primary sensory cortices, unimodal association cortices, multimodal cortices as well as structures several synapses removed from sensory inputs and motor outputs that they might be considered amodal. (p.107)

According to Chatterjee, information is progressively "bleached" of modal content as it is transmitted successively farther from primary sensory areas. He suggests that productive work can be done by framing questions according to several functional-neuroanatomical axes that capture this graded transition from modal to amodal representation—a transition that trades specificity and referential grounding, for flexibility and generative power. Instead of asking whether representations are modal or amodal, Chatterjee suggests we ask what can be done with representations that are modal and amodal to varying degrees. All these synthetic proposals account for the empirical results generated by the modal camp while retaining the strengths of amodal approaches.

While these hybrid proposals maintain some version of the modal-amodal distinction that we criticized above (§4.2), they nevertheless push the field in what we see as a productive direction. Hybrid accounts all share a distrust of a unitary, invariant neural system for conceptual representation, and a commitment to multiple, context-sensitive neural systems. Rather than attempting to characterize the format of cognitive representations, we believe that a productive research program will investigate how multiple, varied cognitive resources support particular cognitive functions in particular contexts. But we suggest that this realization can be pushed one

step further. In the next section, we argue that the study of conceptual processing would be best served by abandoning the modal-amodal distinction entirely, and focusing instead on the *functional role of functionally specialized neural circuits*.

5. The proposal: Functional roles of functional circuits

We have argued that debates about the modality or amodality of mental representations and processes, if assessed at the neural level, are unresolvable. But this does not imply that there is no room to ask how various neural processes, including those also involved in perception and action, contribute to language comprehension and other aspects of higher cognition. In fact, these kinds of questions are essential and productive—as long as the questions don't implicitly assume a principled distinction between modal and amodal. Instead of investigating the format of mental representations or processes—trying to pin down whether they are "modal" or "amodal"—research of the sort conducted under the banner of the Modal Hypothesis should study instead the mechanisms that are functionally responsible for particular components of higher cognitive accomplishments. For instance, a researcher concerned with language could ask: What functional role is played by which parts of the brain (and body) in well-defined aspects of language learning, production, and comprehension? This is qualitatively different from questions about representational format—whether modal, amodal, or multimodal. It's a question not about format but about mechanism, and specifically about shared mechanisms. Questions about mechanism are necessary if we intend to understand how humans are capable of language comprehension and of higher cognition more generally.

5.1 Functional roles and the mechanisms of cognition

If we want to know how the brain is capable of higher cognition, we will naturally need to consider a lot of neural territory. Let's take the example of language once again. Like any complex cognitive behavior, processing language activates a greedy web of cortical and subcortical areas. The research reviewed in Section 3 above showed that some of those brain areas are also implicated in perception and action, and moreover that the activation of these areas is sensitive to linguistic content. We also saw that comprehending language about action and perception interacts with perceiving and acting in systematic ways, so these interactions don't just take place at the neural level—they also affect behavior. As a result, we take it as uncontroversial that neural resources that are responsible for

aspects of perception and action are also active during language comprehension. What remains to be seen is whether their activity is functionally relevant to comprehension, and if so, *how*.

But it is not a stretch to hypothesize—and at present this is merely a hypothesis—that these same regions might indeed play functional roles in both language comprehension, and perception and action. This claim is usually implicit in modal accounts, and for good reason. Parts of the brain that serve particular functions for perception and action seem exceptionally well suited to a range of functional roles in comprehension. For one, they are good candidates for producing subjective experiences; if the mechanisms responsible for perception are also responsible for understanding language about a given percept, then whatever subjective experience one has for the former may also be experienced for the latter. Likewise, if simulated action and perception share neural machinery with real action and perception, then simulated actions and percepts may inherit the mechanisms that perform inferences in response to the genuine article (cf. Grush, 2004), thus supporting inference during language comprehension. And when it gets right down to it, the functional properties of brain structures responsible for identifying objects and events during perception might make those brain structures particularly well suited for representing those same objects and events in their absence, in exactly the way suggested by Barsalou (1999) and others in the modal tradition.

This suggests a mechanistic reframing of the Modal Hypothesis that does not reify the distinction between modal and amodal representations:

> **The Shared Functional Circuit Hypothesis:** Functional neural circuits—assemblages of neurons that coordinate to accomplish a particular function—can play a functional role in superficially dissimilar cognitive tasks.

This *Shared Functional Circuit* approach highlights shared neural machinery between disparate cognitive accomplishments, and focuses on the functional contributions of that machinery. The question of functional role, however, is of a very different character from those usually asked in research undertaken to support the Modal Hypothesis. We're not advocating asking merely whether some piece of brain tissue in V1, MT+, or premotor cortex is activated, for whatever purpose, during a higher cognitive function like language comprehension; this would simply add to the literature showing that those cells are active during behaviors other than perception or action.

Instead, we're arguing that we should ask what that neural activity contributes, functionally, to the behavior in question.

Answering this question will require several angles of attack. What properties of the stimulus, the environment, or the cognizer—their history, their individual characteristics, their mental or physical state—affect the likelihood that this piece of tissue will be involved? What leads to the absence of activity? What, functionally, does the presence (or absence) of activity do to the cognitive process in question? For language comprehension, for example, we might find that when people don't use circuits in MT+, they have less detailed mental representations of described motion, or perhaps they are less able to draw inferences pertaining to the motion of those objects—two quite different functional contributions that are glossed over in debates between modal and amodal camps. Instead of debating the representational format of a neuron or neural system, we ought to try to find out what that cell or system does—the behavioral and cognitive functions that it subserves.

5.2 From perceptual format to perceptual function

It might be attractive to move toward the functional role of shared functional circuits, and away from representational format, while maintaining some sense in which those circuits are perceptual or motor (cf. Gallese & Lakoff, 2005). Suppose we found that some aspect of language comprehension—the representation of described objects, for instance—was functionally supported under certain circumstances by parts of V1. We might be tempted to conclude that this was a case in which systems dedicated to perception and action are being re-used functionally for other cognitive purposes; thus, higher cognition would be "modal." As intuitively appealing as this might be, it amounts to the same line of reasoning that we argued against in Section 4: making format attributions on the basis of shared neural resources, when matters of format dissolve at the neural level. The move to define multifunctional neurons as primarily or intrinsically perceptual or motor is neurally unfeasible.

This isn't to say that the perceptual or motor functions of some particular neural circuit are uninformative. On the contrary, if evidence suggests that some functional circuit plays a role in language comprehension but is *also* functionally implicated in perception or action, then this can provide clues to the functional contribution of this circuit to comprehension. Knowing that MT+ is

responsible for motion perception, for instance, suggests that it may play a similar role in language comprehension—such as representing described motion (Saygin et al., 2010).

There's a delicate balance to be reached here. Researchers need a way to designate certain pieces of brain tissue as being of interest because they have known functions in perception or action, but without problematically having to label them as ontologically "perceptual", "motor", or "modal"—more specifically, without making claims about format, claims that we've argued dissolve at the neural level. With care, we can avoid this difficulty. Instead of broadly targeting perceptual or motor systems, we can instead focus on specific networks of neurons that serve demonstrated functional roles during perception or action. Without committing ourselves to the origin of this specialization, we take it as given that the adult brain is highly specialized, with particular neural populations performing particular computations for particular tasks. This specialization may come about by dint of innate predisposition for a particular type of connectivity, a lifetime of training, or a combination of these. The upshot of this, crucially, is that we have neural circuits specialized for controlling effector-specific motor actions in M1, circuits that process motion in MT+, and so on. That specialization allows us to ask whether a functional circuit that spends a lot of its time computing direction of motion from visual input, for instance, also bears some functional load during language comprehension, either generating motion-related inferences, or preparing for appropriate action, or representing described motion.

So, in place of identifying neurons (or neural systems) as perceptual or motor or amodal, *tout court*, we propose instead to take as the object of neural study *functional circuits*: connected networks of neurons that demonstrably perform stable functions for given cognitive behaviors. Observing that some neural circuit performs a given set of computations during perception does not commit us to classifying it as "perceptual." The focus shifts from identifying format to identifying the computations performed by a neural circuit, and asking whether that functionality is also recruited for other behaviors—and if so, which ones, when, why, and with what consequences.

5.3 Finding functional circuits

How can we find functional circuits of interest? We need to be careful, of course, to avoid purely topographical definitions of circuits, lest we fall into the trap of neo-phrenology. That said, the locations of networks of neurons can inform questions about the functional assemblages in which

they participate; assuming typical development, neurons in V1 are likely to be attuned to input from the eyes, for instance. But only an analysis of the dynamics of a neural circuit can indicate what it does, and for which cognitive operations. In principle, functional circuits may differ substantially in their promiscuity, from true one-trick functional ponies, only ever engaged for a single type of processing in response to a specific type of input, to more widely deployed circuits that contribute to a variety of behaviors—although recently authors have argued that the latter are far more common (e.g., Anderson, 2010). The ones most relevant to the question at hand are those keenly tuned to performing specific functions during perceptual or motor tasks, but which also play a functional role in some aspect of comprehension or other conceptual processing.

Adopting this strategy doesn't lose anything in terms of the broad scientific concerns that drove research on the Modal Hypothesis. And it still allows us to ask the core questions that motivate work on cognition, like how the human brain is able to perform operations like language comprehension, and whether evolutionarily older brain circuitry used for perception or action has been exapted for language use. But it does so in a piecemeal and non-programmatic way, defining circuits by their functions and asking what the (possibly multifarious) contribution(s) are of each functional circuit, within and across behaviors. Liz Bates famously said that language is a "new machine built out of old parts" (Bates, 1992). Focusing on functionally defined brain circuits allows us to test this hypothesis in a detailed way, asking whether these old parts (specialized circuits that do certain things during perceptual or motor tasks) are used for other purposes, and what those purposes might be.

To be clear, we do not mean to suggest that functional circuits are informationally encapsulated, like Fodorian modules (Fodor, 1983); a functional circuit may be informationally integrated with other circuits. Nor do we mean to reify these functional circuits into a singular ontology of non-overlapping functional circuits, a brain's bestiary. A given neuron might be seen to participate in more than one functional circuit, depending on context, input from other cortical circuits, or even just the level of analysis relevant for understanding the cognitive accomplishment at hand. And the organization of neurons into functional circuits may change over time, throughout development or in response to the current situation. In many ways, we are agnostic about the best way to characterize these functional circuits; indeed, this is the very task at hand. We merely want to suggest that for the purposes of understanding how the brain accomplishes some complex cognitive exploit

like language comprehension, neural populations can be fruitfully studied in a mechanistic analysis (Bechtel, 2008) that pinpoints those neural circuits that make functional contributions—that is, those circuits without which the relevant cognitive processes and their observable products would be different[2].

There is good reason to ask whether circuits that perform perceptual or motor functions are also used to perform other cognitive accomplishments. Fortunately, in a framework that seeks the functional roles of functional circuits, we can still make progress on that question without falling into the representational quicksand of labeling particular neurons as modal or amodal. The Shared Functional Circuit approach promises a mechanistic answer, instead of a format answer, to questions about the relation between perception, action, and cognition.

6. Some leading lines of inquiry

If our aim is to determine what functional role (if any) certain known neural circuits play in specific aspects of language comprehension (or any other cognitive function), then this will limit the empirical methods we can conscript for this task. In this section, we outline some design considerations: properties that would allow experimental work to inform the question of shared functional circuitry. We then identify several recent studies that make the most progress in this direction.

6.1 Design features of an experimental approach to the functional role of functional circuits

The Shared Functional Circuits approach makes two demands. The first is to pick out functional circuits—populations of neurons that coordinate to perform a particular function. The second is to demonstrate that those neural circuits make functional contributions to more than one cognitive accomplishment: perhaps motion perception and understanding motion language, or color perception and representing the color of discourse referents. Satisfying the first demand requires that we vary the task and dependent measure, so as to narrow down the function of the neural circuit—for language comprehension, this could be causal inference, online representation of a discourse model, memory for linguistic referents, among many other functions. Determining the functional properties of a neural circuit requires a combination of methodologies, from neuroimaging techniques that can identify location and timecourse, to careful and varied behavioral

manipulations that can identify the circuit's precise functional profile. Moreover, the granularity at which we pick out these functional circuits will depend on the granularity of the questions we're asking. Investigating the representation of the visual appearance of discourse referents will require that we pick out one functional circuit; investigating the representation of the *color* of discourse referents, a more delimited circuit. A large body of research responds to this first demand, much of it surveyed in Section 3.1 above.

Satisfying the second demand is more difficult—and this is has been the source of considerable criticism of the Modal Hypothesis. Much existing behavioral work on language, perception, and action (reviewed in Section 3) uses facilitatory priming of perception or action by language. Well-known examples include sentential priming of picture processing (Zwaan et al., 2002) or manual actions (Glenberg & Kaschak, 2002), among others. Experiments like these may *suggest* the involvement in language comprehension of circuits that perform perceptual or motor functions, but in themselves they can never identify the functional contribution to language comprehension, if any. That is because, as discussed, these results are vulnerable to the Mediating Process Reply: it could very well be the case that some other circuitry performs all the heavy lifting for language comprehension, and that priming is observed only because of spreading activation to the relevant circuits involved in perception or action.

The same is true for behavioral experiments in which the order of the tasks is reversed, with a perceptual or motor task facilitating some subsequent linguistic behavior. For instance, participants might be faster to read and make a judgment about a described action when they have just finished performing that same action. If the sensorimotor behavior reliably activates a given circuit, then finding sensorimotor-task-to-language facilitation would tell us that particular circuit can activate the meanings of particular words or phrases—potentially a novel finding in cases like abstract language (e.g. Wilson & Gibbs, 2007). Such findings, however, do not necessarily license the stronger inference that some aspect of language comprehension *relies* on some part of the neural substrate of the sensorimotor task. Circuits that control the relevant perception or action could be associatively linked to relevant linguistic circuits—again, the Mediating Process Reply—so the facilitatory priming of language does not demonstrate a functional role for a given functional circuit. The same holds, importantly, for TMS studies that demonstrate facilitation of language processing by application of a magnetic field to an area believed to house sensorimotor circuits (Pulvermüller et al., 2005; Willems

et al., 2011). It could again be the case that increased activation of the sensorimotor circuits lead, through association, to activation of other, language-specific circuitry.

For a similar reason, brain imaging alone cannot tell us whether particular circuits perform a given functional role in language comprehension, and if they do, what functional role that might be. The detection of an increased BOLD signal associated with language of a particular type and in a particular brain region does not imply, as discussed above in Section 3, that there is any functionality to the underlying brain activity for the particular cognitive processes under consideration. Activity in visual areas during language comprehension, for instance, could merely reflect downstream activation that *results from*, but does not *contribute to,* the comprehension process. And for that matter, there could well be neural circuitry used for language—but not perception or motor control—that is strategically housed in close proximity to circuits that do have sensorimotor functions. This isn't to say that neuroimaging is not informative; indeed, imaging studies can finger neural circuits as likely candidates for redeployment. Yang and Shu (2012), for instance, applied Granger Causality analysis to fMRI data collected during action verb comprehension. They found evidence of bidirectional causal relations between premotor cortex and the posterior medial temporal gyrus, an area thought to be involved in lexical-semantic processing, while activity in primary motor cortex appeared to be causally dependent on the posterior medial temporal gyrus. These results suggest that premotor cortex, but not primary motor cortex, may play a functional role in language comprehension—but they remain unable to unambiguously demonstrate a functional role, or to specify what that role might be. Similarly, van Elk and colleagues (2011) conducted an EEG study that found activation of motor areas during the processing of action verbs. This motor activation preceded the N400, an ERP component thought to index difficulty with semantic integration, and the motor activation was inversely proportional to the amplitude of the N400. These results suggest that activity in motor areas can facilitate semantic integration, but again cannot rule out the possibility that other brain regions mediate between motor activity and lexical-semantic processing. The current temporal and spatial precision of brain imaging, therefore, does not allow us to exclude the possibility of epiphenomenal activation or the involvement of neighboring but distinct neural circuits, and thus imaging results remain vulnerable to a Mediating Process interpretation.

We've argued that facilitation effects and neuroimaging results are inconclusive in that they could be due either to functional engagement of a neural resource or to epiphenomenal activity attributed to

spreading activation. Although such epiphenomenal activity could incidentally result in facilitation of a subsequent task, the same is less obviously true for interference effects, since certain types of interference effect are thought to reflect competition for the neural resources that play a functional role in a cognitive task (Bergen, 2007). If a cognitive process is *inhibited* because a particular cognitive resource is unavailable, then we *can* infer that the process relies functionally on that resource.

More specifically, the key to inferring functional roles for functional circuits involves *knockout effects*: interference effects in which we observe selectively degraded performance on one task as a result of another task, treatment, or trauma that "knocks out" the ability to use some known neural circuit. The basic logic is to find a way to interfere with a neural circuit that serves a known function for perception or motor control, for instance, and to determine what other functions (if any) are selectively affected—for instance, accessing the meaning of words or sentences, making inferences, responding appropriately to language, etc. Knockout effects afford clearer inferences about the functional overlaps between linguistic and sensorimotor tasks to the extent that they are known to be due to the unavailability of a particular resource, thus forestalling recourse to a Mediating Processes Reply.

In order to make these inferences, however, two things must be established. First, the sensorimotor task has to affect the specific linguistic manipulation being studied, not just language processing across the board; the same basic logic applies here as to demonstrations of dissociations in general. In practice, this can be established by including a control condition with linguistic stimuli not thought to rely on the relevant functional circuit. For instance, one might want to show that knocking out the neural circuitry that detects motion towards the body interferes selectively with processing sentences related to motion towards the body, and should therefore also include sentences that describe motion away from the body. Even stronger, a double dissociation logic could be used, such that knocking out circuits that detect motion towards or away form the body uniquely interfere with processing sentences about language towards or away from the body, respectively.

Second, it has to be established that the interference effect is indeed *interference* on the condition of interest, rather than facilitation of the other experimental condition[s]. To continue with our motion example, for instance, we would need a baseline measure that establishes the reaction time to

process a sentence (or whatever the dependent measure is) absent the specific knockout manipulation. If reactions in the motion-toward condition are slower than in the baseline condition, but reactions in the motion-away condition are not slower than in their own baseline condition, then this licenses the inference that processing of this particular type of sentence is selectively impaired when the motion-relevant neural circuit is knocked out. Comparison to a baseline condition, therefore, is necessary to distinguish interference from facilitation effects, which are vulnerable to Mediating Process Replies.

The canonical class of knockout effects are neuropsychological lesion studies in which trauma to particular parts of the brain is shown to selectively impair specific cognitive deficits. Imagine a study in which lesions to neural circuits known to play a particular role in effector-specific motor control are found to impair comprehension of some aspect of language about actions using those same effectors but not others. On the basis of this study, we could conclude not only that the lesioned circuits also perform a relevant function for language, but also something about the function performed. But knockout studies need not involve permanent damage to neural tissue. The same logic would apply to a hypothetical study that used TMS to disrupt a localized functional circuit, and as a result impaired not only a perceptual or motor function but also some aspect of language comprehension. And it would also apply to behavioral work, as long as it was possible to behaviorally impede the use of a neural circuit with a known sensorimotor function, and as long as knocking out that circuit again impaired some aspect of a concurrent or subsequent language task.

Alas, there are some further wrinkles. The Mediating Process Reply still potentially applies to some knockout results. For instance, interfering with the use of a circuit with a sensorimotor function (by engaging it in an unrelated sensorimotor task or even possibly through TMS) could again produce downstream activation of a mediating circuit, and thereby interfere with that circuit. Thus, simultaneous or subsequent interference on a language task might arise due to *downstream* knock-out of a mediating circuit. For instance, if a sensorimotor task requiring motion perception also selectively interferes with concurrent comprehension of motion language (as in Kaschak et al., 2005), then this may be due to the unavailability of the functional circuits responsible for motion perception, as commonly concluded. Alternatively, however, the Mediating Process Reply raises the possibility that the motion-perception task may activate functional circuits that are responsible for motion perception, which incidentally activate other downstream circuitry that represents the

concepts AWAY or TOWARD, and—since activating a particular cognitive resource can leave it less available for other tasks—this could therefore leave these latter resources less available for language comprehension. Crucially, motion perception and language comprehension would on this account rely functionally on entirely different functional circuits. While this scenario may seem unlikely, it nevertheless reveals that knockout experiments are not unambiguous as to whether or not the "knocked-out" functional circuit is functionally implicated in the secondary linguistic or cognitive task.

There are at least two ways to address this lacuna and militate against a Mediating Process Reply: use an adaptation paradigm, or include a mediating-circuit control task. In an adaptation experiment, one task is repeated or continued until the neural resources subserving the task become less available for a nontrivial "refractory" period *after the task is stopped*. A well-known example is motion adaptation, in which viewing motion in one direction for an extended period of time produces a subsequent bias to perceive ambiguous motion as moving in the opposite direction (Mather et al., 1998). If a language task relies on some of the same functional circuitry as the initial task, then performance on the language task should be selectively impaired during the refractory period.

Adaptation is a clear way to implement a knockout design behaviorally, but how does it evade the Mediating Process Reply? How do we know that interference on a language task is not due to the adaptation of other downstream processes? We reason as follows. As a given circuit starts to habituate to a stimulus—say, as a motion-detecting circuit adapts to a motion stimulus—that circuit begins to fatigue. This is the heart of the adaptation paradigm. But, crucially, this fatigue will dampen any potential downstream activation to other circuits. As the first circuit habituates, therefore, these hypothetical downstream circuits—the target of Mediating Process Replies—should receive less input from the first circuit, and would thus be less likely to habituate. Said otherwise, if the initial adaptation task incidentally activates and thus disrupts some downstream circuit, then the demands on this downstream mediating process should *decrease* as the initial circuitry becomes fatigued and thus less reactive to input. So if the interference on the subsequent language task is due to this downstream mediating process, then the observed interference should actually diminish or even *disappear* with longer adaptation periods. Thus, if an associated linguistic task is inhibited after an adaptation task—and if the inhibition persists or even increases with longer adaptation periods—then we can safely conclude that some of the functional circuitry subserving the original

sensorimotor task is also functionally implicated in the linguistic task. If this reasoning holds, then adaption paradigms can unambiguously demonstrate shared functional circuitry.

The second way for a behavioral knockout experiment to forestall a Mediating Process Reply is to include an mediating-circuit control task. If completing a sensorimotor task interferes with a concurrent or subsequent language task, and if this interference is due, as the Mediating Process Reply would argue, to the downstream activation of some language-specific circuit, then qualitatively comparable interference should arise if we *directly* activate the purported mediating language circuit. Adding a control task that attempts to activate the mediating process can test this possibility. If perceiving motion *away from* or *toward* the participant is found to interfere selectively with their comprehension of language describing motion away or toward, then a control task should attempt to activate the language-specific representations invoked by a Mediating Process Reply—the symbols AWAY and TOWARD, for instance. One way to do this would be to use bimodal stimulus presentation, where the words "away" or "toward" are presented visually while participants respond to sentences presented auditorially. In order to rule out the mediation of language-specific circuitry, we would need to see that the mediating-circuit control task does not produce the same pattern of interference results as the sensorimotor task.

The foregoing discussion identified desiderata for experiments that investigate shared neural resources—shared, for instance, by language, perception, and action—that cannot be explained by recourse to downstream activation. A knockdown argument for shared functional circuitry requires experiments that knock out access to the relevant circuits, and asks how knocking them out affects the relevant behavior. We've also argued that behavioral knockout studies ought to additionally incorporate adaptation or mediating-circuit control tasks in order to deal with potential Mediating Process Replies. Of course, it is unlikely that any such knockout experiment will show a complete inability to understand language, since comprehension is not a unitary phenomenon but the collaborative accomplishment of various functional circuits; interfering with any particular functional circuit should only dampen, not destroy, language comprehension. The focus is not on necessity for language comprehension, but on the specific functional contributions of particular functional circuits. And interference on a language task, modulated selectively by a concurrent or preceding sensorimotor task, and with the appropriate controls in place, is the kind of evidence necessary to uncover the functional circuits responsible for the varied aspects of language use.

6.2 Some current work

There is a small body of recent work that exemplifies the type of knockout experiment we've defended. In contrast to the studies discussed in Section 3, the studies discussed here all used a linguistic task as the dependent variable, looked for (and found) interference in matching conditions, and some also implemented appropriate controls. We highlight them here to suggest both that this type of research is possible, and that early results suggest that neural circuits responsible for perception and action may, indeed, play functional roles in language use.

Meteyard and colleagues (2008) examined the effect of a perceptual task on language processing. Participants looked at random dot kinematograms, in which some dots moved randomly, and some moved coherently either upward or downward. After looking at a kinematograms, subjects made lexical decisions to visually presented words. Compared with a control condition, viewing coherent motion at subjects' perceptual threshold produced significant interference on incongruent motion verbs. Subjects were slower to respond to verbs describing downward motion like 'fall' after viewing upward-moving kinematograms, but they responded equally quickly to verbs describing congruent motion and to control verbs such as 'eat' and 'kick.' Thus, the effect was interference for incongruent motion verbs rather than facilitation for congruent verbs. This is a relatively straightforward dissociative knockout effect, and it exemplifies many of the desiderata from the preceding section. However, the lack of a mediating-circuit control task makes it impossible to rule out the possibility that viewing the kinematograms activated other language-specific circuitry due to downstream spreading activation, which was in turn responsible for the observed decrement in performance.

More recently, Casteel (2011) found interference on language tasks following engagement of the motor system (see also Yee et al., in press). Casteel (2011) collected reading times for short passages describing motor actions. Before and after reading each passage, participants were instructed to imagine performing (Experiment 1) or to actually pantomime (Experiment 2) a motor action that was either a match or a mismatch for the action that would be described in the written passage. Participants took longer to read passages that matched the action compared to mismatch trials and to trials with no action at all. A third experiment showed that this interference effect disappeared when the action was performed only before reading but not after, leading Casteel (2011) to conclude

that it was the maintenance of motor codes in working memory that interfered with language comprehension involving similar movements. This work is highly suggestive, but again would have been strengthened by the addition of a mediating-circuit control task—such as a working memory task that required participants to remember the name of the action.

Perhaps the result that best fulfills all the design desiderata outlined above is the series of studies reported by Glenberg and colleagues (2008). They adapted the motor system with a repetitive task in which participants moved 600 beads either away from or toward themselves, depending on their condition. After completing this bead task, participants had to judge the sensibility of sentences, including some that described motion away or towards. The bead task selectively interfered with performance on the linguistic task: Sensibility judgments were slower when bead motion and the described motion matched, but otherwise unaffected (or sometimes slightly facilitated). In addition, the authors further forestalled a Mediating Process Reply with an mediating-circuit control task in which participants had to make 600 judgments of whether a visual stimulus was the word "away" or "toward," or an anagram of that word (e.g., "towrad"); unlike the selective interference observed in the bead experiments, there was no interaction between the motion in the mediating-circuit control task ("away" or "toward") and the motion described by the sentences in the linguistic task.

In sum, there is gradually accumulating evidence that certain functional circuits make functional contributions to perception and action, and to the comprehension of corresponding language. Many issues remain to be resolved, including the role of timing in determining facilitation/interference effects, and whether the findings previously observed in facilitation results can also be observed in interference paradigms. We hope that future studies can be strengthened by taking into account some of the design suggestions discussed above.

7. Conclusion

Questions of representational format are compelling because they promise to address the workings of higher cognitive functions, like language, including how they emerged over evolution and development. The idea that "modal systems" contribute to higher cognition is a particularly compelling answer, one in which the mind is a "new machine built from old parts" (Bates, 1992), with systems that evolved for perception and action now exapted for use in newer functions like

mental imagery, language, or mathematical reasoning. However, we've argued here that attention to representational *format* cannot live up to this promise—indeed, that the very notion of representational format, as currently used in the debate, is untenable. Instead, we've advocated an alternative approach, one that asks: What neural circuits with known functions, particularly functions in perception and action, are also used for which linguistic (and other higher cognitive) purposes?

The Functional Circuits approach dispenses with issues of format while allowing us to ask about the mechanisms of higher cognition. And to the extent that it turns out that there are functional circuits that play roles in both sensorimotor and higher cognition, the answer resembles the Modal Hypothesis with which we started. The approach championed by modal theorists is probably a productive way to identify candidate circuits for study, since the evolutionarily older circuits involved in perceiving and acting are particularly ripe for re-use (Anderson, 2010). But the Shared Functional Circuits approach is also considerably different from the Modal Hypothesis, in at least two ways.

First, we may discover that the functional circuits reused for language are not limited to ones also specialized for perception or motor control. The early evidence, reviewed above, does seem to suggest a role for such machinery in language comprehension. But language may also recycle entirely different circuits that support functions outside perception and action, like analogy, categorization, generalization, binding, etc. (cf. Goldstone & Barsalou, 1998). A motley crew of functional machinery might well conspire to enable particular aspects of language use; indeed, cultural practices often coordinate diverse cognitive resources to accomplish complex cognitive work (Hutchins, 2010). Mathematics, for instance, may involve the skillful deployment of evolutionarily older circuits that subserve spatial cognition (e.g., Dehaene et al., 1993; Dehaene & Cohen, 2007), but also more evolutionarily-recent abilities like analogical reasoning or conceptual integration (Gentner, 2003; Fauconnier & Turner, 2002; Lakoff & Núñez, 2000). The crucial point is that attending to the functional roles of functional circuits has the potential to explain how, mechanistically, we are able to do things like understand language—regardless of whether those circuits also have roles in perception and action.

Second, the Modal and Functional approaches differ in the kind of questions they ask about this new machine built from old parts. Instead of asking whether language understanding uses modal

systems, we now ask what functional roles in the new machine are performed by which circuits, and what other cognitive capacities those circuits might support (cf. Bergen, 2012, ch. 10). Because we advocate asking about mechanism and function rather than format, the answers to those questions will follow in kind. If we're asking about the role in language of circuits that also contribute to perception or action, the type of answer we get will be of the form: such-and-such a neural circuit plays such-and-such a function in language use. For instance, we might find that circuits that are specialized for motion perception play important roles in key aspects of language comprehension, like accessing or representing meaning, performing inference, or generating the subjective experience of comprehension, among others. Indeed, this way of reframing the debate situates the Modal literature within larger theoretical discussions of the massive redeployment or recycling of neural circuitry (Anderson, 2010; Dehaene & Cohen, 2007).

In arguing against the possibility of distinguishing representational formats at the neural level, our target has not been the notion of representation itself. Others have argued in various ways against the importance of mental representation, especially when static representations overshadow the dynamism and situatedness of thought (e.g., Gibson, 1979; Suchman, 1987; Varela et al., 1991; Brooks, 1991; Thelen & Smith, 1994; Chemero, 2009). We have not engaged with that debate here. We recognize, however, that the Shared Functional Circuits hypothesis shifts attention away from representation, and toward neural dynamics and function. While our argument does not place any pressure on the very notion of mental representation, the Functional Circuits approach may offer a bridge between dynamical systems theorists, connectionists, classical representationalists, enactivists, and various other camps.

The Shared Functional Circuit Hypothesis, therefore, reframes the debate surrounding the Modal Hypothesis so we can ask a different type of question, at once broader and more focused. The answers it invites, framed in terms of the functional roles of functional circuits, can make concrete and constructive contributions to the ongoing science of the mind, contributions that are easily translated into mechanistic models of language use that make testable predictions. This would be a laudable improvement over the current state of affairs, mired in terminological debates. After all, what better contribution to the study of language and higher cognition than to document, without representational prejudice, the sundry neural resources that conspire to generate our greatest cognitive accomplishments?

Acknowledgements

Thanks to Vic Ferreira, Jeff Elman, Seana Coulson, and Rafael Núñez for critical feedback on earlier drafts of this article.

References

Anderson, M. L. (2010). Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences,* 33, 245–313

Aydede, M. (1999). What Makes Perceptual Symbols Perceptual? Peer commentary on Barsalou (1999). *Behavioral and Brain Sciences*, 22, 610-611

Aziz-Zadeh, L., Wilson, S., Rizzolatti, G., & Iacoboni, M. (2006). A comparison of premotor areas activated by action observation and action phrases. *Current Biology*, 16, 1818-1823.

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

Barsalou, L.W. (2008). Grounded cognition. *Annual Review of Psychology, 59,* 617-645

Barsalou, L.W. (2010). Grounded cognition: Past, present, and future. *Topics in Cognitive Science, 2*, 716-724.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning* (245–283). Oxford, UK: Oxford University Press.

Barsalou, L.W., Simmons, W.K., Barbey, A.L., & Wilson, C.D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences,* 7, 84-91.

Bates, E. (1992). Language development. *Current Opinion in Neurobiology,* 2, 180-185.

Bates, E., Wilson, S.M., Saygin, A.P., Dick, F., Sereno, M.I., Knight, R., Dronkers, N. (2003) Voxel-based lesion-symptom mapping. *Nature Neuroscience*, 6 (5), 448-450.

Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.

Bergen, B. (2007). Experimental methods for simulation semantics. In M. Gonzalez-Marquez, I. Mittelberg, S. Coulson, and M.J. Spivey (eds.), *Methods in Cognitive Linguistics:* Ithaca.

Bergen, B. (2012). *Louder Than Words: The New Science of How the Mind Makes Meaning.* New York: Basic Books

Bergen, B. & Chang, N. (To appear). Embodied Construction Grammar. In Thomas Hoffmann and Graeme Trousdale (eds.). *The Oxford Handbook of Construction Grammar*. Oxford.

Bergen, B., & Wheeler, K. (2010). Grammatical aspect and mental simulation. *Brain & Language,* 112, 150-158.

Binder, J.R., & Desai, R. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences,* 15, 527-536.

Bolognini, N., & Ro, T. (2010). Transcranial magnetic stimulation: Disrupting neural activity to alter and assess brain function. *Journal of Neuroscience,* 30, 9647-9650.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence,* 47, 139-159.

Buccino G., Riggio L., Melli G., Binkofski F., Gallese V., & Rizzolatti G. (2005). Listening to action-related sentences modulates the activity of the motor system: A combined TMS and behavioral study. *Cognitive Brain Research,* 24, 355-363.

Carey, S. (2009). *The Origin of Concepts.* New York: Oxford University Press.

Casteel, M.A. (2011). The influence of motor simulations on language comprehension. *Acta Psychologica,* 138, 211-218.

Chatterjee, A. (2010). Disembodying Cognition. *Language and Cognition.* 2, 79-116.

Chemero, A. (2009). *Radical Embodied Cognitive Science.* Cambridge, MA: MIT Press

Connell, L. (2007). Representing object colour in language comprehension. *Cognition,* 102, 476-485.

Dale, R. (2008). The possibility of a pluralist cognitive science. *Journal of Experimental and Theoretical Artificial Intelligence*, *20*, 155-179.

Davies, W.M. (2004). Amodal or perceptual symbol systems: A false dichotomy?. *Behavioral and Brain Sciences*, 27, 162-163.

Dehaene, S., Bossini, S., & Giraux, P. (1993). The mental representation of parity and numerical magnitude. *Journal of Experimental Psychology: General*, 122, 371-396.

Dehaene, S., and Cohen, L. (2007). Cultural Recycling of Cortical Maps. *Neuron,* 56, 384-398.

Dove G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition,* 110, 412–431.

Fauconnier, G. and Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, New York.

Fodor, J. A. (1975). *The language of thought.* New York, NY: Thomas Y. Crowell.

Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology,* 22, 455-479.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (eds.), *Language in mind: Advances in the study of language and thought*, 195–235. Cambridge, MA: MIT Press.

Geschwind, N. & Kaplan, E. (1962). A human cerebral deconnection syndrome. A preliminary report. *Neurology,* 12, 675-685.

Ghazanfar, A.A., & Schroeder, C.E. (2006). Is the neocortex essentially multisensory? *Trends in Cognitive Science*, 10, 278-285.

Gibson, J.J. (1979). *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, 9 (3), 558-565.

Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory & Language,* 43 (3), 379-401.

Glenberg, A. M., Sato, M., & Cattaneo, L. (2008). Use-induced motor plasticity affects the processing of abstract and concrete language. *Current Biology,* 18, R290-R291.

Goldstone, R., & Barsalou, L. (1998). Reuniting perception and conception. *Cognition,*65, 231-265

Grice, H.P. (1962). Some remarks about the senses. In R. J. Butler (Ed.), *Analytical Philosophy, First Series.* Oxford University Press.

Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences,* 27, 377-396.

Hadamard, J. (1954). *The Psychology of Invention in the Mathematical Field.* New York: Dover.

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

Hume, D. (1739/1978). *A treatise of human nature.* Nidditch, P. H., ed. Oxford: Oxford University Press

Hutchins, E. (2008). The role of cultural practices in the emergence of modern human intelligence. *Philosophical Transactions of the Royal Society of London Series B,* 363, 2011-2019.

Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason.* Chicago: University of Chicago Press.

Johnson-Frey, S. H. (2004). The neural bases of complex tool use in humans. *Trends in Cognitive Sciences,* 8, 71–8.

Johnson-Laird, P.N. (1983). *Mental Models.* Cambridge: Harvard University Press.

Kaplan, D., & Bechtel, W. (2011). Dynamical Models: An Alternative or Complement to Mechanistic Explanations. *Topics in Cognitive Science,* 3, 438-444.

Kaschak, M. P., Madden, C. J., Therriault, D. J., Yaxley, R. H., Aveyard, M. E., Blanchard, A. A., & Zwaan, R. A. (2005). Perception of motion affects language processing. *Cognition*, 94, B79–B89.

Kayser C., & Logothetis, N.K. (2007). Do early sensory cortices integrate cross-modal information? *Brain Structure and Function,* 212, 121-132.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.

Kosslyn, S.M. (1980). *Image and Mind*. Cambridge: Harvard University Press.

Kosslyn, S.M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, MA: MIT Press.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.

Lakoff, G., & Nunez, R. (2000). *Where mathematics comes from*. New York: Basic Books.

Lemus, L., Hernández, A., Luna, R., Zainos, A., & Romo, R. (2010). Do Sensory Cortices Process More than One Sensory Modality during Perceptual Judgments? *Neuron*, 67, 335-348.

Locke, J. (1690/1979). *An essay concerning human understanding*. P. H. Nidditch, ed. Oxford: Oxford University Press.

Machery, E. (2007). Concept Empiricism: A Methodological Critique. *Cognition*, 104, 19-46.

MacPherson, F. (2011). Individuating the senses. In: MacPherson, F. (ed.) *The Senses: Classic and Contemporary Philosophical Perspectives*. Oxford University Press, pp. 3-46

Mahon, B.Z., & 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.

Markman, A., & Dietrich, E. (2000). Extending the classical view of representation. *Trends in Cognitive Science,* 4, 470-475.

Marr, D. (1982) *Vision*. New York: W.H. Freeman and Sons.

Masson, M. E. J., Bub, D. N., & Warren, C. M. (2008). Kicking calculators: Contribution of embodied representations to sentence comprehension. *Journal of Memory and Language*, 59, 256-265.

Mather G., Verstraten, F.A.J., & Anstis, S. (1998). *The Motion Aftereffect: A Modern Perspective*. Cambridge, MA: The MIT Press.

Meteyard, L., Rodriguez Cuardrado, S., Bahrami, B. & Vigliocco, G. (2012). Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex,* 48, 788-904.

Meteyard, L., Zokaei, N., Bahrami, B., & Vigliocco, G. (2008). Visual motion interferes with lexical decision on motion words. *Current Biology*, 18, R732-R733.

Marr, D. (1982). *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman and Company.

Matlock, T. (2004). Fictive motion as cognitive simulation. *Memory & Cognition*, 32, 1389-1400.

Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. MIT Press.

Minsky, M. L. (1977). A framework for representing knowledge. In: *The psychology of computer vision*, ed. Ph. H. Winston. McGraw-Hill.

Newell, A., & Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliffs: Prentice-Hall.

Nichols, S., Stich, S., Leslie, A., & Klein, D. (1996). Varieties of Off-Line Simulation. In P. Carruthers & P.K. Smith (Eds.), *Theories of Theories of Mind.* Cambridge: Cambridge University Press

Ochipa, C., Rothi, L. J. G., & Heilman, K. M. (1989). Ideational apraxia: A deficit in tool selection and use. *Annals of Neurology*, 25, 190 193.

Oliveri, M., Finocchiaro, C., Shapiro, K., Gangitano, M., Caramazza, A., & Pascual-Leone, A. (2004). All talk and no action: A transcranial magnetic stimulation study of motor cortex activation during action word production. *Journal of Cognitive Neuroscience,* 16, 374-381.

Paivio, A. (2007). *Mind and its Evolution: A Dual Coding Theoretical Approach.* Mahwah, NJ: Erlbaum.

Pezzulo, G., Barsalou, L.W., Cangelosi, A., Fischer, M.A., McRae, K., Spivey, M. (2011). The mechanics of embodiment: A dialogue on embodiment and computational modeling. *Frontiers in Cognition*, *2*(5), 1-21.

Prinz, J. J. (2002). Furnishing the mind: Concepts and their perceptual basis. Cambridge, MA: MIT Press.

Pulvermüller, F., Hauk, O., Nikulin, V., & Ilmoniemi, R.J. (2005). Functional links between motor and language systems. *European Journal of Neuroscience,* 21, 793-797.

Pylyshyn, Z.W. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences,* 25, 157-238.

Pylyshyn, Z.W. (2003). *Seeing and Visualizing: It's Not What You Think.* Cambridge, MA: MIT Press

Richardson, D., Spivey, M., Barsalou, L., & McRae, K. (2003). Spatial representations activated during real-time comprehension of verbs. *Cognitive Science,* 27, 767-780.

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience,* 27, 169-192.

Rizzolatti, G., & Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. *Nature Reviews Neuroscience,* 11, 264-274.

Saygin, A.P., McCullough, S., Alac, M., & Emmorey, K. (2010). Modulation of BOLD response in motion sensitive lateral temporal cortex by real and fictive motion sentences. *Journal of Cognitive Neuroscience*, 22 (11), 2480-90.

Simon, H. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society,* 106, 467-82

Sirigu, A., Cohen, L., Duhamel, J.R., Pillon, B., Dubois, B., & Agid, Y. (1995). A selective impairment of hand posture for object utilization in apraxia. *Cortex,* 31, 41-55.

Sober, E. (1993). Preface. In: *Conceptual Issues in Evolutionary Biology*, ed. E. Sober, ix-xviii. MIT Press.

Stanfield, R.A., & Zwaan, R.A., (2001). The effect of implied orientation derived from verbal context on picture recognition. *Psychological Science,* 12, 153-156.

Tettamanti M., Buccino G., Saccuman M.C., Gallese V., Danna M., Scifo P., Fazio F., Rizzolatti G., Cappa S.F., & Perani D. (2005). Listening to Action-related Sentences Activates Fronto-parietal Motor Circuits. *Journal of Cognitive Neuroscience, 17, 273-8*.

Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action.* Cambridge, Mass.: MIT Press.

Thomas, N.J.T. (forthcoming). Mental Imagery. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2012 ed.). Retrieved from http://plato.stanford.edu/archives/win2012/entries/mental-imagery/

Varela, F. Thompson, E., & Rosch, E. (1992). *The embodied mind: Human cognition and experience.* Cambridge, Mass.: MIT Press.

Willems, R.M., Labruna, L., D'Esposito, M., Ivry, R., & Casasanto, D. (2011). A Functional Role for the Motor System in Language Understanding: Evidence From Theta-Burst Transcranial Magnetic Stimulation. *Psychological Science,* 22, 849-854.

Wilson, N.L., & Gibbs, R.W., Jr. (2007). Real and imagined body movement primes metaphor comprehension. *Cognitive Science,* 31, 721-731.

Yang, J., & Shu, H. (2012). The role of the premotor cortex and the primary motor cortex in action verb comprehension: Evidence from Granger causality analysis. *Brain Research Bulletin*, 88, 460-466.

Yaxley, R.H., Zwaan, R.A., (2007). Simulating visibility during language comprehension. *Cognition,* 105, 229-236.

Yee, E., Chrysikou, E., Hoffman, E., & Thompson-Schill, S.L. (in press). Manual experience shapes object representation. *Psychological Science.*

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

Zwaan, R.A., & Taylor, L.J. (2006). Seeing, acting, understanding: Motor resonance in language comprehension. *Journal of Experimental Psychology: General*, 135, 1–11.

Notes

[1] There is an active philosophical debate about how best to individuate our senses into distinct modalities, but neural architecture on its own is seldom a candidate criterion (Grice, 1962; MacPherson, 2011).

[2] At this point, the reader with allegiances to Dynamical Systems accounts of cognition may balk at our appeal to functionally specialized neural circuits. According to Dynamical Systems theorists, cognitive processing is inherently *interaction-dominant* (Spivey, 2007), characterized more by interactions *between* components than by the isolated contributions of individual components. Is the

Shared Functional Circuits approach committed to a component-dominant architecture? We don't think so. For starters, the appeal to mechanism is not antithetical to dynamical systems approaches (Kaplan and Bechtel, 2011). Additionally, the Shared Functional Circuits approach does not require that the brain be "nearly decomposable" (cf. Simon, 1968), only that certain neural populations have a stable and recyclable function—said otherwise, that they have a reliable effect on the trajectory of neural activity through state space. The Shared Functional Circuits approach, therefore, is complementary to Dynamical Systems approaches that admit "soft modularity" (Spivey, 2007, p. 123).