Embodiment, Consciousness, and the Massively Representational Mind

Robert D. Rupert
University of Colorado, Boulder

ABSTRACT. In this paper, I claim that extant empirical data do not support a radically embodied understanding of the mind but, instead, suggest (along with a variety of other results) a massively representational view. According to this massively representational view, the brain is rife with representations that possess overlapping and redundant content, and many of these represent other mental representations or derive their content from them. Moreover, many behavioral phenomena associated with attention and consciousness are best explained by the coordinated activity of units with redundant content. I finish by arguing that this massively representational picture challenges the reliability of a priori theorizing about consciousness.

Proponents of the embodied approach to cognitive science (Barsalou et al. 2003, Gibbs 2006; Glenberg 1997; Lakoff and Johnson 1999) claim that embodiment-related empirical results support sweeping, negative conclusions: that there are no amodal symbols; that there are no arbitrary symbols; that functionalism is false; that computationalism is false; and more. At least a grain of truth, though one best paired with a grain of salt, grounds each of these claims. Elsewhere I argue for this moderate diagnosis (Rupert 2006; 2009, chapter 11), mostly by criticizing the arguments

for radical embodiment-related conclusions. In the present essay, I aim to be more positive, in two ways.

First, I argue that results to do with embodiment help to reveal something important about cognitive processing: in conjunction with a wealth of further work in cognitive science, they suggest a massively representational theory of human cognition. Embodiment-related experimental results seem obviously to show that bodily posture, movements, and emotional states affect cognitive processing. I contend that they have such effects precisely because the experimental stimuli in question activate bodily *representations*—i.e., either (a) representations *of* bodily states or (b) representations the physical realizers of which are in distinctively bodily (that is, sensory or motor) cortex, regardless of whether they represent bodily states or, instead, some other aspect of reality. These bodily representations contribute to the production of behavior alongside a plethora of other representations—many with the same content as the bodily ones, but many of which are realized in cortex that is not particularly bodily. On this view, although the ultimate account of our cognitive processing may be computationalist in some respects, it must take into account the widespread effects of a buzz of representations, many with redundant content.[1, 2]

Second, I am keen to sort out the relation between the embodied movement and consciousness. Limitations of space preclude anything close to a full exploration of these issues, but in a brief, closing section, I argue that certain aspects of the mind's massively representational, associative workings deflate the dualist-seeming import of common intuitions about conscious experience (Chalmers 1996; cf. Dennett 1988, 1991). In some ways, this conclusion should come as no surprise; many advocates for the embodied view have thought the approach sheds important light on the nature of consciousness. Nevertheless, I offer only Pyrrhic victory to the fans of embodiment. A significant proportion of the philosophical work on embodiment invokes a first-person, phenomenological perspective, rife with uncashed appeals to conscious experience (see, for example, various essays in Stewart et al. 2011). In contrast, the line I press tends to undermine the reliability of such reports; their role is neither that of foundational phenomena to be explained nor direct conduit to the theoretical nature of conscious states. Thus, even if subjects can report reliably that they are in one conscious state rather than another, their doing so serves only as one kind of data, among many other kinds, that an overarching theory of cognition must account for. More importantly, claims about the abstract nature or properties of such states (that they have intrinsic qualitative characters, for instance) should be treated as suspect, even if subjects can reliably report that they're in one conscious state (for example, seeing red) rather than another (for example, seeing blue). I contend that this deflationary treatment of the theoretical value of first-person reports on conscious states falls out of our best-supported theories of cognition—theories that draw heavily on the embodied literature.

I. REPRESENTATIONS: NUMEROSITY, REDUNDANCY, CONTENT, AND COLLABORATION

In this section, I introduce four theses concerning mental representation and the human mind-brain and describe empirical results that illustrate and support the four theses. To clear the ground, however, I remark first on the liberal way in which I understand the idea of a mental representation.

On the view of mental representations adopted here, a mental representation corresponds to, or is about, an individual, property, or kind, correspondence to which helps to explain (successfully and according to fairly strict standards of scientific explanation) intelligent behavior. One might best treat a mental representation as an element in a psychological model that can be paired with another element in the model, such that (1) the former is mapped onto a state internal to the target cognitive system, (2) the former bears, according to the model, some kind of privileged relation to the latter (e.g., the latter reliably causes the activation of the former or the former tracks the presence of the latter), and (3) the relation referred to in (2) plays a role in successful applications of the model to target (psychological) systems (that is, in the explanation of the intelligent behavior of the target system). A simple example will help to bring these remarks down to earth. Imagine I run into a friend at the grocery, who introduces me to a third party, call him 'Charlie'. A week later, I pass Charlie in the square, and I say 'hello'. Assume that some unit was distinctively active when I first met Charlie and that its renewed activity guided my recognition of Charlie a week later. In this case there are (at least) three elements in the model—those mapped onto (1) the "cognitive" unit, (2) Charlie, and (3) my greeting behavior in the presence of Charlie (a week after initially meeting him). The later activation of the unit that was distinctively active when meeting Charlie accounts for the *explanandum*: my uttering a greeting in the presence of Charlie a week after meeting him. So the co-presence of the unit and Charlie at the initial time of the meeting, and perhaps an ongoing privileged relation (however precisely one's theory of representation spells that out), plays an explanatory role with respect to the modeling of the *explanandum*.

I will take this to be a respectable notion of representation. Much has been written about the relation between this sort of notion of representation, on the one hand, and conceptual analyses of representation and content, on the other (Fodor 2008; McDowell 1994; Ramsey 2007; Rowlands 2006). Without wading fully into those waters, it should be clear from the literature on naturalistic semantics that the relation in question is not mere causation; extant proposals are much more demanding than that (Dretske 1988; Fodor 1990; Rupert 1999; Ryder 2004).[3] In my view, this question ultimately concerns empirical matters; if the relevant aspects of intelligent behavior can be accounted for by invoking a thinner notion of representation than some authors have had in mind, then so be it. Such a result would show that the thin sort of representation is all the representation there is, regardless of

what it's called. One can think of my goal here as an attempt to see how much causal-explanatory mileage can be wrung from only (or at least primarily) a relatively thin sort of representational-relation. On, then, to the four theses.

> 1. *The massive number of representations (MNR):* The mind-brain contains an enormous number of basic mental representations, a number far exceeding what is suggested by introspection or by standard philosophical theories.

What number of mental representations does introspection or do standard philosophical theories suggest? I take as stalking horse the view that a subject's stock of atomic representations is roughly equal in size to her natural-language vocabulary. Perhaps the introspective experience of one's own internal voice while thinking recommends such a view; more to the point, a grand, roughly Fregean tradition bequeaths this view upon us: belief-state attributions in natural language track, or even are somehow constitutive of, mental content. In contrast, if *MNR* is correct, the sheer number of "orthographically distinct" mental representations is orders of magnitude larger than our natural-language vocabulary. (I limit *MNR*'s scope to *atomic* mental representations; given the operation of combinatorial functions, a vast number of compound representations can be constructed from virtually any stock of atoms—a result that would trivialize *MNR*.)

> 2. *The massive redundancy of mental representations (RED):* For any given subject S, for very many representational contents C, S possesses many more than one vehicle with C as its content.

In other words, for whatever property, individual, or kind a subject sometimes represents, she has lots of ways of representing it—where having "different ways of representing *C*" amounts to nothing more than having distinct, psychological-level units that represent *C*.

> 3. *Internal-externalism about content (IEC):* Representational vehicles can have other vehicles or other vehicles' contents as their own representational contents—in virtue of the same sorts of relations that determine externalist content.

When I speak of representational content, I have in mind externalistic content of the sort operative in the example involving Charlie, above. Whatever else such content is, it is not intrinsic to vehicles; moreover, something's having externalist content is not tantamount to the vehicle's being associated with a sense or idea that then determines the vehicle's referential or truth-conditional properties and to which the subject has direct or *a priori* access (Putnam 1975). Ultimately, I suspect that our best theories of cognition will deal in externalist content only, partly because such theories, if they can be made to work, better satisfy naturalistic scruples. I do not argue for such a view here. As suggested above, I mean to see only how much can be productively explained by invoking only this form of content.[4]

Thesis *IEC* applies externalist thinking internalistically, focusing only on what is essential to the externalist-style determination of content; it is content determined at least partly by factors external to the vehicle or to any idea or sense asso-

ciated with that vehicle, but *not* necessarily content determined by factors external to the agent or body in which that vehicle appears. When contrasting externalist notions of content with internalist ones, philosophers often deploy the twins test (Putnam 1975): if molecule-for-molecule duplicates differ in mental content, this shows that at least one of the molecular duplicates (and more plausibly, both) has at least one mental vehicle (or representation, or concept) with content determined by its relation to elements beyond her bodily boundary. This test does in fact specify a sufficient condition for externalist content. Moreover, reflection on the result of applications of the twins test can aid our theorizing, perhaps indispensably, by helping us to home in on the conditions that determine external reference. Nevertheless, when one internal mental representation bears, for example, the germane causal relation to another internal vehicle, the first has the second as its "externalist" content, even though a twins test, treated as a necessary and sufficient condition on the possession of externalist content, might not deliver this result; for, depending on the nature of the germane causal relation, duplication may bring said relation in its train when that relation holds among internal units.[5]

Here I mean to leave open various possibilities with regard to the way in which one internal vehicle might have internal-externalist content. For example, vehicle *r* might simply represent vehicle *s*, qua vehicle (or perhaps represent one of *s*'s non-representational properties). Or, it's possible for a vehicle to have internal-externalist content by deference to the externalist content of another vehicle: in this way, *r* might simply inherit the externalist content of another vehicle *s* (which externalist content could be itself inherited from yet another vehicle, and so on); that one vehicle inherits its content from a second would plausibly be determined by the functional role of the former vehicle, as fixed by functional architecture, functional decomposition, or evolutionary history. Or, *r* might inherit *s*'s externalist content partly by representing *s*, *qua* vehicle; *r*'s content might thus have two simultaneous aspects: *r* might be about *s and* about whatever *s* is about.

> 4. *Thesis of coordination (CO):* The co-contribution of redundant representations to the production of behavior explains many phenomena normally associated with attention and consciousness.

A greater or smaller number of redundant representations might causally contribute to the production of an individual instance of behavior. I contend that variation in the number involved—perhaps as a proportion of all representations that are active—explains variance in such behavioral measures as error rates and reaction times, and, in some cases, the subject's very ability to exhibit reliably the behavior in question. Moreover, an increased number of co-contributors explains the increased task-success (reduced reaction times and error rates) that we would normally attribute to someone's *paying attention* or *being consciously focused on the task at hand*.[6]

How might redundant representations "team up" to contribute to more effective (faster, less error-prone) behavior? They might do so by distinct causal pathways (e.g., via distinct motor assemblies), but just as likely, they contribute via

partially distinct pathways, creating stronger input to a single cortical area, higher levels of activity in which enhance performance. The particular apportioning of contributions among the various redundant representations—and the effectiveness of their co-activation—might depend on historical connections. In some cases, the internal causal relations that ground content also create and sustain "recruitment" relations. Let's say that vehicle r comes to have s's content as its own as the result of being causally linked to s during performance; r is activated by s and when this happens, it enhances the subject's efficiency in handling the external object, property, or individual represented by s, so that r's being correlated with what s represents helps to explain (causally, and by fairly strict standards of scientific explanation) why r's being activated by s enhances performance. The formation and continued manifestation of this relation can reinforce associative links that allow one vehicle (r or s) to recruit the other, enhancing performance as a result (which enhancement persists even after content has been determined). Vehicle r has s (or s's content) as its content in virtue of r's having been recruited for control. The recruitment was behaviorally successful. So, there's a correlation between *(a)* the two vehicles' sameness of content and *(b)* their co-contribution to enhanced behavior (Dretske 1988; Jackson and Pettit 1988).

The following example illustrates the four theses. In the 1980s, Patricia Goldman-Rakic (1987) did groundbreaking work on the role of feedback-based processes during the execution of tasks that require working-memory—involving visuospatial information, in particular. She investigated "the mechanism(s) by which symbolic representations are accessed and held on line to guide behavior in the absence of (or in spite of) discriminative stimuli in the outside world" (ibid., 380). Much of the relevant data comes for studies on monkeys, but based on arguments by homology and by analogy to human dysfunction, Goldman-Rakic suggests that the following sort of process is operative in the human case as well as that of monkeys. Early visual processing activates a collection of feature-representations that themselves activate a parietal representation, which causes a corresponding representation to be "loaded" into working memory in frontal cortex. The maintenance of a feedback loop—representations in parietal cortex activating representations in frontal cortex, which reactivate the parietal representations, and so on—accounts for successful performance on tasks requiring that information be held in working memory.

Here we find more representations than might have been expected—a parietal representation in addition to the one in prefrontal cortex—and so support for *MNR* (further connections between the parieto-frontal circuit and motor control and the laying down of long-term memories suggest even more representations—ibid.). The process also illustrates *RED*: the parietal and prefrontal units represent the same information about the environment. Whether this case supports *IEC* depends on somewhat murky empirical matters, for example, whether the prefrontal units act as dynamically re-assignable pointers gaining their external content on the fly, by "demonstrating" back to the parietal units (thereby borrowing

the content of the parietal units, a clear instance of *IEC*), or whether parietal units merely causally mediate the connection between standing types of prefrontal units and the external properties or structures they represent. Lastly, the activation of redundant representations helps to explain successful performance on attention-related tasks: disrupt the feedback loop, thereby reducing the number of redundant representations that are active, and performance drops off.

This research does not emerge from a particularly embodied perspective (although note that one kind of representation that can be maintained in this feedback loop is information about bodily position and orientation—ibid., 381); Goldman-Rakic hopes primarily to explain a range of cognitive deficits associated with damage to dorso-lateral prefrontal cortex, and thus to explain the contribution of that cortex to cognitive functioning.

II. EVIDENCE OF THE MASSIVELY REPRESENTATIONAL MIND

In this section, I summarize a variety of results—some more to do with embodiment, some less so—that illustrate and support the four theses, the collection of which I refer to hereafter as *MRIC*. No single case, embodiment-related or otherwise, establishes *MRIC*; rather, it's the aggregate of cases that impresses me. Moreover, the extent to which evidence for the presence of a given kind of mechanism or process supports *MRIC* depends on the extent to which that mechanism or process plays a role throughout the cognitive system, a matter that remains uncertain in many cases. Caveats issued, however, the massively representational view captures a recurring pattern, one manifest in model after model, successful research program after successful research program.

Some further comments about methodology and the dialectic are in order. My central claim is that variation in the number of co-contributing redundant representations accounts for a significant amount of variance in behavioral measures associated with intelligence, attention, and conscious awareness of a task. Boiled down even further, the message is this: *number of redundant representations contributing to the production of behavior* should be considered a theoretically important property in cognitive science. Thus, *RED* and *CO* form the core of this paper's thesis. That being said, it's difficult to see how, in actual systems, *RED* could be true, if *MNR* were not. In contrast, *IEC* may seem entirely optional. Nonetheless, *IEC* helps us to make sense of some important cognitive processes (see the discussion of subitizing and FINST models below) and also bears on larger philosophical questions, to do with consciousness, for instance.

A. VISUAL ENHANCEMENT OF TOUCH

Consider first the phenomenon of the visual enhancement of touch (Kennett et al. 2001; Serino and Haggard 2009). When subjects visually perceive a body part that

is, at the same time, subject to tactile stimulation, subjects make more accurate judgments about tactile stimulation, and they make some judgments more quickly. Magnification of the visually presented body part enhances this effect.

Here's a more detailed description of the typical experiment. The experimenter stimulates the surface of the subject's skin with a mechanical device that has two adjustable protruding pointers. When the two pointers are set close enough together (below the threshold of discrimination), subjects judge that only one pointer is pressing against their skin. Experimenters do not allow subjects to see the tactile stimulation itself. The body part—the arm, for example—rests inside an enclosure, so subjects can't simply look to see how many pointers protrude from the instrument. Enhancement is achieved by allowing the subject to see the relevant body part through a window (which can also be fitted with a magnifying lens, in place of simple see-through plastic film). When the subjects see the relevant body part, their discrimination thresholds become smaller; subjects can tell that two points are simultaneously touching their arms, even when those points are so close together that they would normally be treated as a single tactile stimulus. (Note, too, that visual objects present in the location of the arm do not enhance discriminations. Thus, the effect is not merely due to spatial attention paid to the location of the arm. Rather, it would seem that activation of the additional bodily representation accounts for enhanced tactile sensitivity.)

How is VET connected to the massively representational view? Here is a rundown.

MNR: The proposed explanation posits interaction between (visually activated) representations in multisensory areas and representations in primary somatosensory cortex.

RED: These multiple representations are of the same area of bodily surface: the tactile representations have the same content as at least part of the visually activated representation.[7]

IEC: Consider the magnification effect. It's plausible that the content of a cluster of otherwise co-representing somatosensory neurons is dynamically altered by the content of visually activated representations (Kennett et al. 2001, 1190). The response profiles become more finely tuned—some of the neurons in the cluster coming to refer to a more specific area of the bodily surface than they normally refer to, and others coming to refer to a different, more specific area than they normally refer to—in a way that could plausibly be interpreted as the result of content inheritance: multisensory units command the somatosensory units to "represent this degree of detail," thus making the content of the somatosensory units at least partly derived from the content of the multisensory units. Moreover, one might think that the multisensory units represent the somatosensory units when issuing their command ("you lot; represent this"), or perhaps the somatosensory units represent the multisensory ones when borrowing their content ("I'm to borrow my content from *those* neurons").

CO: Enhanced tactile sensitivity results specifically from the cooperation of multiple redundant representations—representations of the body part in multisen-

sory areas and representations in primary somatosensory cortex. This has the effect one might expect when a subject pays careful attention. Paying more attention normally enhances one's discriminative abilities: when one looks more carefully, one sees more detail, and, as a result, makes finer discriminations.

B. FORWARD MODELS, EFFERENCE COPY, AND PREDICTIVE CODING

Consider the use of forward models in human cognition (Grush 1997, 2004). The essential idea is as follows. When one part of cortex sends a command—say a motor command to reach or to move one's eyes—a copy of that signal, known as efference copy, is sent elsewhere in the brain for processing. Such processing consists in a comparison of the simulated results of that command to a goal state. This process can allow for a new, corrective motor command to be sent if the simulated result does not match the goal state.

Whereas one might have thought that the brain-mind issues a single command to move a limb, the efference copy increases that number—as does the use of representations in simulating the outcome, as does the use of updated motor commands in response to mismatch (*MNR*). There is no doubt, too, that inclusion of forward models in our account of cognitive processing introduces representational redundancy: the effective motor command as well as the efference copy represent the same action (*RED*). With regard to *IEC*, matters are slightly less clear. Much depends on the details of the architecture, but one plausible way of filling in the details supports *IEC*. A forward model might be indifferent to which limb is involved but contain a pointer back to the source of the command: "Whichever arm you commanded to reach for your hat will not go high enough." (There is some evidence of the existence of functionally specialized neural areas that represent actions in a limb-independent way [Butterfill and Sinigaglia 2012; Cattaneo et al. 2010; Chaminade et al. 2005; Rijntjes et al., 1999; Williams and Gribble 2012]; granted, this work is not concerned with forward models, or activity of cerebellum, which is thought to be the site of forward models [Grush 2004, 385], but it is at least suggestive.) The relation to *CO* seems clearer. The matching of two representations with the same content—the representation generated by the forward model and representation of the goal state—helps to explain attentive, hitch-free behavior.

Beyond its role in the deployment of forward models, efference copy's use is likely to be widespread, as part of the use of predictive coding, for instance (Clark, forthcoming). It would appear that a wide variety of cognitive processes depend on the brain's own prediction of incoming signals, such that when the predicted signals match those coming in (sensory observations of the result of an action, say) the prediction signals and efferent signals "match" and cancel each other out, whereas a mismatch leads to error correction and learning. This use of predictive coding supports *MRIC*: (*MNR*) every prediction generates additional representations; (*RED*) the predictive signal can, on one reasonable interpretation, be taken to represent the state of the body or world represented by the incoming signal, and,

when all goes well, the contents of these two match; (*IEC*) the top-down signal refers to a state of the world or body by referring to the *type* of expected signal; we understand why the corrective mechanism works if we understand it as representing, even if only indirectly, where the limb was expected to be (I expect the arm to be *there*, where the content of 'there' is a function of the predictive signal's referring to the relevant type of input signal, that is, the type of input signal that refers to the limb's being at the location in question); and (*CO*) fluid movement and categorization is more likely to occur when the contents of the predictive and input signals match.

C. VISUAL INDEXES

Turn now to models of early visual processing, in particular, what are known as FINST models (Pylyshyn 2000, 2001). According to the FINST-based approach, early visual processing exploits approximately four pointers each of which can latch onto an object in the world and track it through various of the object's changes, providing a locus of attention for downstream processes. Subjects are able, for example, to track four cued shapes reliably while the shapes move among a field of qualitatively identical distractors that are also moving (Pylyshyn 2001, 142). In a static display, only the properties of units that are cued determine the nature of a visual task aimed at those units (say, judging whether the units are of a certain color), even if further units are present and have properties that, under other circumstances, would act as distractors, altering reaction times or error rates on the task at hand (ibid., 140).

There's also evidence that our ability to subitize rests on the deployment of FINST pointers (Trick and Pylyshyn 1994). Humans enumerate small numbers of objects (four or fewer, typically) quickly and automatically, without much of a per-item increase in reaction times; this is known as 'subitizing'. In contrast, enumeration of five or more items comes with a significant per-item increase in reaction times. On Trick and Pylyshyn's view, indexes first get attached to objects, then downstream processes register the number of indexes active. When, however, the number of items to be enumerated exceeds the number of indexes, additional processes (e.g., the marking of already counted objects) contribute to the production of an answer; as a result, the process can no longer proceed purely preattentively, and response-times increase.

How does such work support *MRIC*? When subjects enumerate, they issue a judgment, which is clearly about the objects represented by the pointers. And in the context of motor control, Pylyshyn remarks, "Indeed, one of the assumptions of indexing theory is that only indexed objects can be the targets of motor commands, including the command to move the gaze to a particular object" (Pylyshyn 2000, 204). Thus, FINST models introduce a new layer of representation, with redundant content, no less, which supports both *MNR* and *RED*. Moreover, the downstream process that produces a subitizing judgment represents the number of objects by representing the number of active indexes and, thus, provides evidence for *IEC*.

Lastly, although this is more speculation than data-driven inference, one might well think that activity directed at objects is more efficient when the motor control processes represent the same number of items as is represented by the collection of active FINST pointers, in keeping with *CO*.

D. OTHERS

In this subsection, I run very briefly through a number of suggestive examples, merely gesturing at ways in which this work might support *MRIC*.

1. Dual-Process Theories

The dual-stream view of vision (Milner and Goodale 1995) has been widely discussed and is supported by a wealth of empirical studies. According to this view, the extraction of features that grounds object categorization, as well as verbal labeling, takes place largely in the ventral stream, proceeding from V1 through V4 and V5/MT to inferotemporal cortex. In contrast, the vision-based construction of an egocentric map that guides fluent bodily action occurs primarily in the dorsal stream, with information flowing from V1 to extrastriate and ultimately posterior parietal cortex. To the extent that both streams represent the same objects or properties in the environment (perhaps in differing formats, so as to allow for different kinds of processes to make use of them), their existence supports *RED*.

2. Concepts

Many readers will be familiar with some aspects of the debate about concepts, which became especially lively in the 1970s with the discovery of family-resemblance structures (Rosch and Mervis 1975). In the decades that followed, the discussion of concepts blossomed (Smith and Medin 1981); among the contending views were the classical view, the prototype view, the stereotype view, and the theory-theory, among others. Recently, Edouard Machery (2009) has argued that this sort of debate is misguided, that a multiplicity of structures serves the purposes to which we tend to think concepts are put—categorization, for example.

On the basis of such multiplicity, Machery argues that concepts are not a natural kind and thus for a form of eliminativism about concepts. Although I'm sympathetic to his eliminativism (about concepts, not about mental representations), such a commitment is beside the point in the present context. Of relevance here is that Machery's picture strongly suggests a mother lode of mental representations. In place of what one might have considered *the* concept COW, there are feature lists, definitions, modally specific templates, normative ideals, and so on.

3. Motion, Time, and Intentions

Now abbreviating even more heavily, I direct the reader's attention to a few further cases. Consider that (a) three different systems detect and represent motion (Lu and Sperling 2001), (b) a variety of different neural mechanisms track temporal intervals (Eagleman 2008, 134), and (c) at least two processes represent the same motor commands or motor intentions (Lau et al. 2007).

4. Representation-Packed Cortex

Consider now the possibility that what is sometimes treated as strength of activation in a given area of cortex might be better understood in terms of the number of redundant representations active in that area of cortex (and similarly for subcortical areas). I use a sample experimental result to illustrate this approach. Moutoussis and Zeki (2002) present subjects with two distinct colored images, one to each eye, that, from the standpoint of conscious report, cancel each other out (subjects report seeing only a homogenous colored field). Nevertheless, fMRI scans indicate that stimuli are being processed by the same areas of the brain that would be active were the stimuli to rise to consciousness. They used color-canceling images of houses and faces, and even under the canceling conditions, there was significant activity in the parahippocampal place area (PPA) and fusiform face area (FFA), respectively. This suggests a picture of cortex densely packed with representational units. The FFA contains many clusters of neurons that, when activated, represent a given face. Under subliminal (in this case, color-canceling) conditions, a smaller number of such clusters become active. Noncanceling, supraliminal conditions activate a greater number of such units. Plausibly, then, the behavior associated with consciousness (e.g., ability to report accurately and verbally on the stimulus) is a function of the number of redundant representations active in the portion of cortex in question: the more face representations active in FFA, for instance, the more likely the subject is to report verbally that she's seeing a face.

III. OBJECTIONS AND REPLIES

In this section, I take up potential objections and respond in ways that should clarify the massively representational view. The responses develop two primary themes. First, I acknowledge that factors other than the sheer number of redundant representations (or the proportion of these relative to the total number of representations active) contribute causally to the production of attentive behavior. I'm satisfied with this result, so long as variation in the number of co-contributing redundant representations accounts for a significant amount of variance in the behavioral data. Second, some of my responses emphasize the appearance of multiple, redundant, co-contributing representations within a single "module" or within what we're inclined to think of as a single internal process; the moral here is that *MRIC* applies even in cases in which it might appear that only a module's single output matters.

OBJECTION #1: MERE MULTIPLE REALIZATION

A familiar enough picture speaks against *RED*. On this view, what I'm treating as redundant representations are, instead, various realizers of the same mental representation; they are not different representations, for representations are individu-

ated by their content alone (and psychological explanation is purely intentional—Fodor 2008). If this view is correct, there probably isn't a massive number of representations after all; for a given subject, the number of equivalence classes of realizers of representational contents might not be too far off of the deliverances of introspection or standard philosophical theories.[8] Whereas I take various empirical results to establish a massive number of distinct representations, often with redundant content, the realization-based strategy opts for a purely intentional principle of individuation for mental representations, in which case, *RED* amounts to the more modest claim that, for any given mental representation, the subject is likely to have many realizers of it or merely has many psychologically equivalent vehicles carrying that representation's content.

Response #1. If redundancy plays the causal-explanatory role claimed by *CO*, the realization-based approach fails decisively. I claim that variation in the sheer number of vehicles with a given representational content plays a causal-explanatory role in the production of certain forms of behavior. A wholly intentional approach to individuation, however, does not accommodate this causal-explanatory contribution; "S occurrently represents C" or "S's mind-brain currently contains as activated representation of C" is true equally of someone who has a single active realizer with content C as it is of a subject that has twenty-five active realizers all of which represent C. Partly because the number of active vehicles does explanatory work in cognitive science, I take being a particular vehicle to be a psychological-level construct—unlike mere realizers, which appear only at some lower level than, or as part of an explanatory enterprise distinct from, psychology.

OBJECTION #2. INTERACTION BETWEEN REPRESENTATIONS WITH DIFFERENT CONTENTS

Generally speaking one might worry that *MRIC* makes no allowance for the presence of *countervailing or conflicting representations*. It makes quite a difference whether twenty representations of *food* are active and nothing else (this is unrealistic, but for the sake of illustration, it will suffice) or twenty representations of *food* and, at the same time, a representation of *poison*. Yet, if we were to focus solely on amount of representational redundancy, it's not clear why the prediction of behavior would differ much—the same amount of food-representing is going on in both cases (and it would be easy enough to emend the objection were we to focus on proportions: imagine that in the first scenario, one instance of PURPLE is active instead of one instance of POISON).

Response #2. A *ceteris paribus* clause does some work in this context: other things being equal, variation in number of representationally redundant units explains a significant amount of the variance in the behavior associated with attention and consciousness. Hold fixed the appearance of a single POISON and the probability of eating the food may still vary with the number of active FOOD-vehicles.

The reader might find this unsatisfactory, for two reasons. First, there is a general concern about the un-informativeness of *ceteris paribus* clauses. At the very

least, we should like to be told something more concrete about the ways in which other factors can mitigate the effects of variations in the number of representationally redundant units. Second, on at least one reasonable understanding of what it is to hold things equal, even a *ceteris paribus* clause can't rescue *CO*. One might think, for example, that, holding fixed the activation of POISON, it makes no difference how many representations of food are active: the content of POISON alone drives behavior in a way that is, for the most part, independent of the content of other active units. (Exception: The content of NOT, as in NOT POISON, may drastically change POISON's causal contribution to behavior, but, again, it would do so in virtue of its content, irrespective of the number of other representations active.)

I acknowledge the allure of a purely content-oriented view, at least as applied to cases such as the poison case. And in the end, I may be forced to concede the existence of phenomena best explained by a purely intentional psychology. All the same, I suspect that the range of such cases is much smaller than *a priori* reflection suggests. It's empirically plausible that cases in which the activation of a representation of POISON prevents a hungry person from eating food involve the activity of a great many vehicles with the content *poison*. Such multiple activations may well be part of the fundamental mechanism of winner-takes-all competitions.

Consider, too, a way in which *MRIC* might contribute to our understanding of cognitive processing even in cases in which the proximate cause of behavior would seem to be amenable to purely intentional explanation. Assume for the moment that such cases involve so-called central processing—a single stream of representations which occupies a position in the architecture such that a single instance of, say, POISON can control behavior regardless of the number of FOOD-vehicles active (because, for instance, a single production rule says "if poison, then, whether food or not food, do not eat"). Fine, but whence the representations that make their way into the central workspace? A standard answer appeals to functionally specialized mechanisms, modules, or neural areas. I do not here wish to endorse a radically modular architecture (Fodor 1983, 2000). Nevertheless, to the extent that relatively isolated, specialized processes provide "final decisions" to central processing, we might ask what happens within those modules. Within such relatively isolated systems, *MRIC* may well apply, determining which outputs emerge from such isolated streams.

What reason is there to think that the number of redundant representations active helps to explain the output of the specialized process? One thought concerns the grain of individuation of vehicles. Discussions of the neural location of representational activity typically advert to the amount of activation in a given neural area, in some cases in order to explain the extent to which a stimulus can control behavior in the way we expect it to when subjects are conscious of it (Moutoussis and Zeki 2002). One plausible gloss of such discussions, suggested above, treats cortex as itself being saturated with distinct representational vehicles and treats variations in strength of cortical activity as being a function of the number of

representations active in an area densely packed with representational vehicles. If this is the correct picture, then variations in, for example, reportability may well be a function, indirectly, of the number of co-representing vehicles active within the local area of cortex.

Notice, too, that *MRIC* might be fruitfully applied in the preceding fashion independent of any commitment to modules or relatively isolated subsystems. Many processes proceed relatively unhindered, at least for short periods of time, regardless of whether they take place within modules or isolated cognitive subsystems. Thus, local competition between densely packed representational vehicles, such that highly active redundancies win out, may constitute a fundamental mechanism of the mind.

OBJECTION #3. EXPERTISE AND AUTOMATIZATION

The acquisition of expertise is normally thought to involve progressive automatization of the relevant thought processes and progressively less and less effortful, conscious thought. This phenomenon might be thought inconsistent with *MRIC*.

Response #3. This objection concerns matters murky. Most importantly, one has to be clear about the *explanandum*, for novices and experts exhibit different forms of attentive behavior. Novices might exhibit painstaking verbal report, though not the sort of attentive behavior associated with intelligence (because they often do not solve the problems set before them). Experts, just the contrary. Both phenomena need explaining, and *MRIC* might shed light on both, but in different ways.

Take the experts first. One plausible account of the changes involved in automatization adverts to increased articulation and proliferation of representational vehicles. As the budding expert becomes more proficient, her behavior is guided increasingly by a greater number of developing representations, operative at the "subconscious" level and perhaps increasingly removed from processes that issue in verbal reports. In contrast, the novice's reports on her reasoning processes may be driven by a rich set of redundant representations, but only within the stream of activity that produces verbal report—which activity might itself be a kind of proficiency associated with attentiveness—even if the novice's problem-solving is not.

V. EMBODIED PROCESSING AND REASONING ABOUT CONSCIOUSNESS

Some authors claim that the embodied view of cognition interweaves deeply, either in its motivations or its implications, with our understanding of consciousness (Gibbs 2006; Stewart et al. 2011; Wheeler 2005). In closing, I suggest a way in which the embodied-cum-massively-representational-view might bear on questions to do with consciousness, but perhaps not in the way embodiment-theorists might have hoped.

Assume that embodied representations contribute to the production of virtually all intelligent acts—where being an embodied representation is being a representation of the state or activity of some part of the body (cf. Gallagher 2005; Prinz 2008)—and that among such embodied representations there are almost always some that do not directly control introspective reports. Assume further that embodied representations are especially likely to be active, in relatively large numbers, when behavior is smooth, fluid, or well executed (and these are especially likely to be embodied representations that do not directly control introspective reports). Assume finally that throughout the preceding sections I've been building up a theory of consciousness—despite all of my roundabout talk concerning *behavior associated with* attention and conscious awareness. As this alternative narrative would have it, I propose in some way to identify phenomenal properties with gradations in the number of redundant representations (perhaps as a proportion of all active representations) that control (or are poised to control) the sort of behavior we associate with attention or consciousness—at least when that behavior is executed well. An interesting result follows: Parts of virtually every conscious state are embodied representations some of which do not directly control introspective report, for example, reports on our cognitive processes, on the contents of our own cognitive states, or on the theoretically important properties of those cognitive processes or states (including properties of the vehicles involved). On this view, the primary materials of *a priori* philosophical theorizing about consciousness—introspective verbal reports or the states that generate them—have no privileged status in the theory of consciousness; rather, it's easy enough to see how they might misdirect philosophers to such conclusions as that conscious experiences have intrinsic qualitative characters for which the natural sciences cannot account (Chalmers 1996).

On what precisely does such skepticism rest? Begin by asking what *does* produce verbal report on conscious experience. I do not intend to offer a theory of what produces verbal report, but we have no reason to think that the processes controlling verbal report have any particular access to the processes in which many embodied representations participate. I emphasize such processes because many of them—the processes that might take place within a system that controls motor behavior by exploiting forward models, for example—seem incredibly unlikely to be processes that one might discover much about via the sort of introspection that produces verbal report.

To my mind, the question of intrinsic qualities stands at the very center of debates over the metaphysics of consciousness (Dennett 1988). And although it may be a straightforward matter for states that control verbal report to track the conscious states, and thus reiterate their content (think *IEC*), it is nevertheless a far cry from there to an accurate verbal report on the theoretical properties of the states the content of which is so borrowed by report-generating processes or of the processes in which the states so reported on participate. Here is where skepticism about the claims of dualists should creep in.[9]

To be clear, I do not challenge the claim that verbal reports carry information (in the dry, technical sense) about the states that produce those reports; that claim

is trivially true. Neither am I challenging the claim that the contents of a subject's verbal reports about conscious states are sometimes satisfied by some of the subject's internal states—particularly some of the states that we associate with sensory experience; perhaps subjects generally get it right when they say "I'm seeing red now." This is all consistent with the thought that theoretical claims about conscious states are quite insecure. On my view, such reports are often the product of theory-generating mechanisms (including ones that construct models of the target processes).[10]

My essential concern, then, which space prevents me from developing in any details, rests on the relation between the stream of (perhaps redundant) representations that produces verbal behavior and the behavior-controlling suites of representations that themselves either constitute experience or constitute proper subparts of conscious experience. Assume that the representations that produce verbal reports have, in most cases, internal-externalist content. They represent other vehicles and, in many cases, borrow the content of those vehicles by doing so. The report-controlling vehicles may be able to ape the content of, say sensory states, and assuming the activation of the report-controlling vehicles reliably tracks types of sensory vehicles, this would account for accurate report (if such there be) of some of our own conscious experiences. Theorizing about the nature of those sensory or conscious experiences is a whole other matter. We can try to use our capacity for formulating models, equations, and such to account for the states that appear, but that kind of theorizing doesn't presuppose direct access to the natures (e.g., intrinsic qualitative characters) of the objects so theorized about.

Lastly, take the processes in which the bodily representations (that do not directly produce introspective reports) participate. That these cooperate to produce behavior, that different representations sometimes track the same property, and so on may be entirely opaque to the report-producing mechanisms and can, at best, be inferred by the theory-generating system on the basis of loads of empirical observation. In other words, I would never have been inclined to produce the immediately preceding sentence had I not been exposed to an enormous amount of experimental results and processed them using the standard model-building mechanisms that drive theorizing about everything from astronomy to particle physics to the motives of others.

It seems plausible enough from, say, an evolutionary perspective that our theory-generating mechanisms would sometimes produce intuitions of unmediated access to truth—better to facilitate quick, decisive action in a hostile environment. Nevertheless, it runs against strong intuitions, of various sorts (for example, Searle's intuition—which seems to be operative in the famous Chinese Room argument—that if someone were to describe what constitutes our own processes of understanding, we would sure enough recognize it as the correct description). Some such intuitions issue from the embodied camp itself, at least insofar as embodied theorists are inclined to rely on phenomenological reports to bolster their position. For my part, I'm inclined to let the empirical work, including work on embodiment, play a trumping role. So far as I can tell, that work calls into question introspective

access to the nature of conscious states, to the way in which their content is determined, and to the workings of cognitive processes in which some of the components of first-order conscious states participate. Thus, embodied theorists can celebrate victory of a sort: a study of embodied representations and processes reveals much about the mind. Nevertheless, such revelation does not flow from *a priori* observations about, or direct introspective access to the nature of, embodied conscious experience.[11]

ACKNOWLEDGMENTS

NOTES

1. Our ultimate account of cognitive processing might take these effects into account by emphasizing associative relations; see Anderson 2007, 37–38; Buckner 2011; Rupert 2009, 9.2, for discussion of complications connected to the contrasts between association and computation and association and cognition.

2. My emphasis on representation contrasts starkly with the dismissive approach prevalent in the literature on embodied cognition, either in the form of an outright denial that the mind traffics in mental representations (Thelen and Smith 1994), a marginalization of mental representations (Chemero 2009), or a radical reconception of them (Brooks 1999). Space does not allow for detailed arguments in support of my representationalist assumption, but my intuitions here have the flavor of an empirical bet. I cannot yet see how a human subject might, for instance, write a book about embodied cognition without deploying loads of mental representations; to read hundreds of articles and books, to retain many of the details and bring them to bear on the production of the manuscript seems, thus far, to defy even a sketchy nonrepresentational explanation. That being said, I might find the pro-representation position compelling partly because I take a liberal view of mental representations (see below). Thus, the disagreement noted here may be a matter of mere semantics (Keijzer 1998).

3. *Contra* Ramsey (2007), it's not *mere causal mediation* either. It's a highly flexible kind of causal mediation that controls the subject's behavior in complex ways that frequently are quite disconnected from the immediate stimulus being tracked or acted upon. Moreover, this behavior often amounts to intelligent interaction with *the very thing represented*, a fact that makes the privileged causal relation identified by extant naturalistic theories seem more like representation than Ramsey recognizes.

4. One complication: in at least some cases, structural properties of complex vehicles might determine representational content (Churchland 1998; Cummins 1996). These cases raise special problems that will not be addressed in detail here. I do not think the appearance of iconic or otherwise holistic representations affects the conclusions of my final section. Internal structure may play, for example, a content-determining role, but I see no reason to think a vehicle's having externalist-

content-determining internal structure itself should ground accurate verbal reports about, or abstract theorizing in connection with, such matters as whether some mental or cognitive states have intrinsic qualitative characters.

It may be worth adding that there is a stipulative aspect to discussions of content. For instance, where I am tempted to talk of associative links between representations with externalistically determined content, someone else might be inclined to talk of a causal-role aspect of content itself. The label doesn't matter to me, so long as we're clear on what role such a construct can and does play (cf. Block 1986).

5. The content of a pointer of the sort commonly employed by computer scientists (Ballard et al. 1997; Newell and Simon 1997/1976, 86–87; Trick and Pylyshyn 1994, 86) would seem trivially to be externalist and *IEC* therefore obviously on the right track, given the widespread use of pointers in the modeling of cognitive processes. I suspect that philosophers tend to ignore the possibility of internal-externalist content because of a commitment to the distinction between what are often labeled 'personal' and 'subpersonal' levels or states. It is common to attribute mental states and abilities to what appears, from a grammatical perspective, to be a whole person; we typically say, "John saw a bear," not "John's visual cortex saw a bear." If the whole person is identified with the entire organism, and we take a whole person to have an immediate grasp on the contents of everything in that whole person's mind, then my rejection of the twins test might seem misguided. All internal vehicles are inside the organism, and so, it might be thought, the person has access to all of the contents in question. This response does not, however, take the teeth out of *IEC*. There must be some relation by which (so-called) subpersonal-level states determine the subject's (supposedly) personal-level states. This relation partly involves principled relations between the contents of subpersonal states and the contents of personal-level states. Thus, if there are content-determining relations among so-called subpersonal states, this is likely to "bubble up": contents of personal-level states will be at least partly a function of contents of the subpersonal-level states, which will itself be partly determined by causal and architectural relations that hold among subpersonal-level states. In this regard, I am primarily concerned—particularly in connection with the later discussion of consciousness—that the personal-level subject has only limited introspective access to the facts about (a) subpersonal contents, (b) causal or architectural relations among subpersonal-level states that determine their contents, and (c) the relation by which the facts of kind (a) and (b) determine contents of personal-level states. In my view, because the subject lacks such introspective access (where introspective access is identified, at a first pass, with the ability to report verbally), we should take a skeptical attitude toward the fruits of our *a priori* theorizing about consciousness.

6. Christopher Mole (2010) and Jesse Prinz (2000, 2008, 2011) have proposed theories roughly in this vicinity, and thus a few comparative comments are in order. First, note that the view on offer here is neither a theory of attention nor a theory of consciousness and so is not, technically, concerned with the same subject matter as either Mole's theory of attention or Prinz's theory of consciousness. All the same, my view is meant to account for some of the same phenomena as their views are meant to account for. In contrast to Mole's view, then, I make no appeal to the personal- or agent-level understanding of a task, which is central to Mole's theory of attention; moreover, I take the measure of the contribution of redundant representations to be modulated by the use of resources that the reflective subject wouldn't naturally think are relevant to the task at hand (they might simply be other areas of cortex that are co-active with the representations in question but that represent something entirely different and thus "distract" the subject or disrupt her performance). In contrast to Prinz's attention-based theory of consciousness, Prinz's requirement that intermediate-level representations be attended to in order that they be components of conscious states limits our explanatory framework unnecessarily, as does Prinz's emphasis on those representations' availability to working memory. That being said, I should acknowledge that my thinking on this issue developed partly as a reaction to Prinz's as well as to Tamar Gendler's emphasis on what she calls 'aliefs' (2008) and Bryce Huebner's (2008, forthcoming) discussion of interactions among subpersonal representations. My view is inspired also by work on consciousness from the 1990s, including Daniel Dennett's popularity-contest view of the mechanisms of verbal report on conscious states (1991) and William Calvin's spreading chorus view (1996).

7. I write "at least part," because I want to leave open the possibility that the visual or multimodal representation is individuated holistically, as, for instance, a single image (Cummins 1996; see note 4, above); in which case, one part of the single visually activated representational structure would

represent the same thing represented by the relevant representational unit in somatosensory cortex; regardless of whether the latter is itself a representational structure with holistically determined content, it will likely map onto a smaller portion of the arm than the entire visually activated unit.

8. Then again, it might be. The view on which this objection is based allows for the appearance of atomic representations that cannot be expressed by single terms in natural language. The present objection challenges *RED*, in the first instance, not *MNR*.

9. The observation driving the preceding note, which I think is correct, calls into question even the accurate reporting of the contents of our cognitive states. Processes that produce verbal report deal in a vocabulary—a way of chunking information—that suits a certain purpose: interpersonal communication. I suspect that many of the vehicles in the motor and perceptual systems represent incredibly fine-grained properties—perhaps expressible only in mathematical formalism—and, as a matter of contingent fact, the processes that generate verbal report do not have a sufficiently fine-grained vocabulary to facilitate reliable tracking of the activation of these representations or of the occurrence of states that activations of these representations partly constitute.

10. Chalmers (1996, 203–204) distinguishes first-order phenomenal judgments and the corresponding sensory registrations of the environment, on the one hand, from second-order awareness of or reflective judgments about experiences, on the other. I claim that *a priori* theorizing about consciousness consists in a kind of third order processing; we construct models of the relation between first-order sensory experiences and second-order registrations of them as well as the relation of both of them to the world and to their contents. Humans might excel at the reporting of some of their own sensory states, and they might do it via simple tracking-based registration (i.e., by covariation in the activations of sensory states with the content, e.g., *red-in-front-now* and the state that controls the verbal report "I see red"); yet, humans might nevertheless be flailing about when they construct theories *a priori* of the processes themselves by which they track first-order sensory experiences of the properties of the vehicles that participate in such processes, etc. Moreover, I suspect those *a priori* theory-constructing operations draw on our standard theory-constructing tools, one side effect of the application of which is to produce more or less conviction (often felt in the form of confidence about intuitions) concerning the products of those theory-building processes.

11. Throughout this section, I have proceeded as if I had earlier been offering a theory of consciousness. I myself am inclined toward eliminativism about consciousness (of the sort that Machery [2009] advocates in connection with concepts). Although the resulting prose numbs the mind, this section's discussion can be recast entirely in terms of the natures of different kinds of cognitive or psychological states and the reliability of various methods of discovering those natures, absent any commitment to a genuine kind or property *being conscious* or *being phenomenal* that is the subject matter of what Chalmers (1996) calls the 'hard problem' of consciousness.

REFERENCES

Anderson, J. R. 2007. *How Can the Human Mind Occur in the Physical Universe?* New York: Oxford University Press.

Ballard, D., M. Hayhoe, P. Pook, and R. Rao. 1997. "Deictic Codes for the Embodiment of Cognition." *Behavioral and Brain Sciences* 20: 723–42.

Barsalou, L., W. K. Simmons, A. Barbey, and C. Wilson. 2003. "Grounding Conceptual Knowledge in Modality-Specific Systems." *Trends in Cognitive Sciences* 7: 84–91.

Block, N. 1986. "Advertisement for a Semantics for Psychology." In *Midwest Studies in Philosophy*: *Studies in the Philosophy of Mind*, vol. 10. Ed. P. French, T. Uehling, and H. Wettstein. Minneapolis: University of Minnesota Press. 615–78.

Brooks, R. 1999. *Cambrian Intelligence: The Early History of the New AI*. Cambridge, MA: MIT Press.

Buckner, C. 2011. "Two Approaches to the Distinction between Cognition and 'Mere Association.'" *International Journal of Comparative Psychology* 24: xx–xx.

Butterfill, S., and C. Sinigaglia. 2012. "Intention and Motor Representation in Purposive Action." *Philosophy and Phenomenological Research*, advanced access.

Calvin, W. 1996. *The Cerebral Code: Thinking a Thought in the Mosaics of the Mind*. Cambridge, MA: MIT Press.

Cattaneo, L., M. Sandrini, and J. Schwarzbach. 2010. "State-Dependent TMS Reveals a Hierarchical Representation of Observed Acts in the Temporal, Parietal, and Premotor Cortices." *Cerebral Cortex* 20: 2252–58.

Chalmers, D. J. 1996. *The Conscious Mind: In Search of a Fundamental Theory.* Oxford: Oxford University Press.

Chaminade, T., A. Meltzoff, and J. Decety. 2005. "An fMRI Study of Imitation: Action Representation and Body Schema." Neuropsychologia 43: 115–27.

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

Churchland, P. M. 1998. "Conceptual Similarity across Sensory and Neural Diversity: The Fodor/Lepore Challenge Answered." *Journal of Philosophy* 95: 5–32.

Clark, A. Forthcoming. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences*.

Cummins, R. 1996. *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.

Dennett, D. 1988. "Quining Qualia." In *Consciousness in Modern Science*. Ed. A. Marcel and E. Bisiach. Oxford: Oxford University Press. 42–77.

Dennett, D. 1991. *Consciousness Explained*. Boston, MA: Little, Brown and Company.

Dretske, F. 1988. *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.

Eagleman, D. M. 2008. "Human Time Perception and Its Illusions." *Current Opinion in Neurobiology* 18: 131–36.

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

Fodor, J. A. 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

Fodor, J. A. 2000. *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.

Fodor, J. A. 2008. *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Gallagher, S. 2005. *How the Body Shapes the Mind*. New York: Oxford University Press.

Gendler, T. 2008. "Alief and Belief." *Journal of Philosophy* 105: 634–63.

Gibbs, R. 2006. *Embodiment and Cognitive Science*. Cambridge: Cambridge University Press.

Glenberg, A. 1997. "What Memory Is For." *Behavioral and Brain Sciences* 20: 1–19.

Goldman-Rakic, P. 1987. "Circuitry of Primate Prefrontal Cortex and Regulation of Behavior by Representational Memory." In *Handbook of Physiology*, vol. 5. Ed. F. Plum and V. Mountcastle. Bethesda, MD: American Physiological Society. 373–417.

Grush, R. 1997. "The Architecture of Representation." *Philosophical Psychology* 10: 5–23.

Grush, R. 2004. "The Emulation Theory of Representation: Motor Control, Imagery, and Perception." *Behavioral and Brain Sciences* 27: 377–96.

Huebner, B. 2008. "Do you see what I see? An investigation of an argument against collective representation. *Philosophical Psychology*.

Huebner, B. forthcoming. *Macrocognition: A Theory of Distributed Monds and Collective Intentionality*. Oxford: Oxford University Press.

Jackson, F., and P. Pettit. 1988. "Functionalism and Broad Content." *Mind* 97: 381–400.

Keijzer, F. A. 1998. "Doing without Representations Which Specify What to Do." *Philosophical Psychology* 11(3): 269–302.

Kennett, S., M. Taylor-Clarke, and P. Haggard. 2001. "Noninformative Vision Improves the Spatial Resolution of Touch in Humans." *Current Biology* 11: 1188–91.

Lakoff, G., and M. Johnson. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.

Lau, H. C., R. D. Rogers, and R. E. Passingham. 2007. "Manipulating the Experienced Onset of Intention after Action Execution." *Journal of Cognitive Neuroscience* 19(1): 1–10.

Lu, Z., and G. Sperling. 2001. "Three-Systems Theory of Human Visual Motion Perception: Review and Update." *Journal of the Optical Society of America* 18(9): 2331–70.

Machery, E. 2009. *Doing without Concepts*. New York: Oxford University Press.

McDowell, J. 1994. "The Content of Perceptual Experience." *Philosophical Quarterly* 44(175): 190–205.

Milner, A. D., and M. A. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

Mole, C. 2010. *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford: Oxford University Press.

Moutoussis, K., and S. Zeki. 2002. "The Relationship between Cortical Activation and Perception Investigated with Invisible Stimuli." *Proceedings of the National Academy of Sciences* 99(14): 9527–32.

Newell, A., and H. Simon. 1997. "Computer Science as Empirical Inquiry: Symbols and Search." In *Mind Design II: Philosophy, Psychology, and Artificial Intelligence*,. Ed. J. Haugeland. Cambridge, MA: MIT Press. 81–110. Reprinted from the *Communication of the Association for Computing Machinery* 19 (March 1976): 113–26.

Prinz, J. J. 2000. "A Neurofunctional Theory of Visual Consciousness." *Consciousness and Cognition* 9: 243–59.

Prinz, J. J. 2008. "Is Consciousness Embodied?" In *Cambridge Handbook of Situated Cognition*,. Ed. P. Robbins and M. Aydede. Cambridge: Cambridge University Press. 419–36.

Prinz, J. J. 2011. "Is Attention Necessary and Sufficient for Consciousness?" In *Attention: Philosophical and Psychological Essays.* Ed. C. Mole, D. Smithies, and W. Wu. Oxford: Oxford University Press. 174–203.

Putnam, H. 1975. "The Meaning of 'Meaning'." In *Mind, Language, and Reality: Philosophical Papers*, vol. 2. Ed. H. Putnam. Cambridge: Cambridge University Press. 215–71.

Pylyshyn, Z. W. 2000. "Situating Vision in the World." *Trends in Cognitive* Sciences 4(5): 197–207.

Pylyshyn, Z. W. 2001. "Visual Indexes, Preconceptual Objects, and Situated Vision." *Cognition* 80: 127–58.

Ramsey, W. 2007. *Representation Reconsidered*. Cambridge: Cambridge University Press.

Rijntjes, M., C. Dettmers, C. Büchel, S. Kiebel, R. Frackowiak, and C. Weiller. 1999. "A Blueprint for Movement: Functional and Anatomical Representations in the Human Motor System." *Journal of Neuroscience* 19(18): 8043–48.

Rosch, E., and C. Mervis. 1975. "Family Resemblances: Studies in the Internal Structure of Categories." *Cognitive Psychology* 7: 573–605.

Rowlands, M. 2006. *Body Language: Representation in Action*. Cambridge, MA: MIT Press.

Rupert, R. D. 1999. "The Best Test Theory of Extension: First Principle(s)." *Mind & Language* 14: 321–55.

Rupert, R. D. 2006. Review of *Embodiment and Cognitive Science* by Raymond Gibbs, in *Notre Dame Philosophical Reviews*, August 20, 2006

Rupert, R. D. 2008. "Frege's Puzzle and Frege Cases: Defending a Quasi-syntactic Solution." *Cognitive Systems Research* 9: 76–91.

Rupert, R. D. 2009. *Cognitive Systems and the Extended Mind*. Oxford: Oxford University Press.

Ryder, D. 2004. "SINBAD Neurosemantics: A Theory of Mental Representation." *Mind & Language* 19(2): 211–40.

Serino, A., and P. Haggard. 2010. "Touch and the Body." *Neuroscience and Biobehavioral Reviews* 34: 224–36.

Smith, E., and D. Medin. 1981. *Categories and Concepts*. Cambridge, MA: Harvard University Press.

Stewart, J., O. Gapenne, and E. A. Di Paolo. 2011. *Enaction: Towards a New Paradigm for Cognitive Science*. Cambridge, MA: MIT Press.

Thelen, E., and L. Smith. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.

Trick, L., and Z. Pylyshyn. 1994. "Why Are Small and Large Numbers Enumerated Differently? A Limited-Capacity Preattentive Stage in Vision." *Psychological Review* 101(1): 80–102.

Wheeler, M. 2005. *Reconstructing the Cognitive World: The Next Step*. Cambridge, MA: MIT Press.

Williams, A., and P. L. Gribble. 2012. "Observed Effector-Independent Motor Learning by Observing." *Journal of Neurophysiology* 107: 1564–70.