2

ASSESSING RADICAL EMBODIMENT

Pierre Jacob

Introduction

Embodied cognition, embedded cognition, enactivism, situated cognition, grounded cognition, and the extended mind are all views of human cognition, human cognitive processes, and the human mind that reject one or another aspect of the Cartesian picture of the mind that survives after Cartesian ontological dualism's replacement by ontological physicalism (cf. Rowlands, 2010). This chapter is devoted to embodied cognition (or embodiment for short). More precisely, it is devoted to what I call *radical embodiment*.

Embodiment can be controversial or uncontroversial. It is uncontroversial that the performance of any human cognitive process causally depends on the possession of a number of physiological and bodily systems or organs such as the digestive, respiratory, and cardiovascular systems. It is also uncontroversial both that human thoughts and beliefs about human anatomy are the output of human cognitive processes and that the truth-conditions of thoughts and beliefs about human anatomy are constituted by facts about human anatomy. But this does not make the relevant cognitive processes embodied. Furthermore, the claim that human visual, auditory, tactile, or olfactory perception is embodied is also uncontroversial to the extent that each kind of perceptual experience noncontroversially depends on the detailed structure of the human eye, ear, skin, or nasal cavity, all of which are noncontroversially parts of the human body. Nor is it controversial that basic human actions (e.g., manual reaching and grasping of objects or locomotion) are embodied to the extent that they involve human bodily movements that noncontroversially depend on the human anatomy.

As Prinz (2008) and Alsmith and de Vignemont (2012) have noted, what is controversial is whether and to what extent either the *possession* and *use* or the *representation* of nonneural bodily parts matters to the possession of human concepts, the entertaining of human thoughts and the execution of higher cognitive tasks

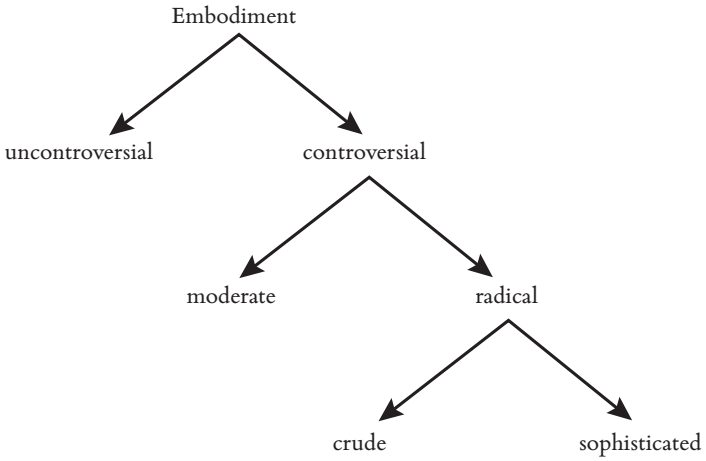

FIGURE 2.1 The varieties of embodiment.

(such as reasoning). Controversial embodiment can be either *moderate* or *radical* (cf. Figure 2.1).

This chapter involves four sections. Because I want to focus on radical embodiment and because radical embodiment is best construed as the rejection of what I will call "Cartesian materialism" (i.e., a Cartesian view of the mind that survives the demise of Cartesian ontological dualism), I will devote the first section to Cartesian materialism. Furthermore, since "Cartesian materialism" sounds like an oxymoron, I will spell out in some details the three basic tenets of Cartesian materialism: the computational-representational approach to the mind, the Fodorian trichotomy, and neurocentrism. In the second section devoted to moderate embodiment, I will argue that moderate embodiment is not incompatible with Cartesian materialism. As I will argue in the third section, the main goal shared by all versions of radical embodiment is the rejection of neurocentrism. I will further draw a distinction between a crude and a sophisticated version of radical embodiment. While advocates of the crude version of radical embodiment are mainly concerned with the further rejection of the computational approach to the mind, advocates of sophisticated radical embodiment are mainly concerned with the elimination of the Fodorian trichotomy. Finally, I shall examine the challenge faced by advocates of radical embodiment who want both to reject neurocentrism and to keep radical embodiment as a genuine alternative to the extended mind thesis (advocated by Clark and Chalmers, 1998). To the extent that Cartesian materialism is intricately linked to the views expressed by Jerry Fodor over the years, this chapter also turns out to a large extent to be an appraisal of Fodor's views.

Cartesian materialism

Several recent influential philosophers of mind (e.g., Dennett, 1991; Haugeland, 1995; Hornsby, 1986; Hacking, 1998; McDowell, 1994; Putnam, 1994) have urged

that some fundamental and objectionable features of the Cartesian picture of the mind can survive and have survived the rejection of Cartesian ontological dualism and the acceptance of ontological physicalism. Thus, Dennett (1991: 107) and Putnam (1994: 488) have coined the expressions "Cartesian materialism" and "Cartesianism cum materialism," respectively. Cartesian materialism can usefully be construed as the conjunction of three fundamental theses: (i) the computational-representational approach to the mind, (ii) the Fodorian trichotomy, and (iii) neurocentrism.

The computational-representational approach to the mind

The computational-representational approach to the mind has been vigorously advocated by Fodor (1987, 1994, 1998). It can in turn be seen as the conjunction of the five following theses (cf. Jacob, 1997 and Horst, 2009):

(i) Mental processes are computational processes.
(ii) Computational processes take mental symbols or representations as input and output.
(iii) The contents or meanings of complex symbols systematically depend on the contents of their constituents and syntactic rules of combination.
(iv) Mental symbols are bearers of underived intentionality.
(v) Psychological explanation is both nomic and intentional, that is, it subsumes psychological events under psychological law-like generalizations that appeal to the contents of an agent's psychological states.

On this picture of the mind, the second *representationalist* assumption is entailed by the first *computational* assumption, in the sense that mental processes could not be fully computational processes unless there existed mental representations that could serve as input and output to mental computations.[1]

What made the computational representational approach to the mind attractive in the first place (cf. Fodor, 1975, 1994, 1998) was its promise to solve an outstanding puzzle inherited from ontological Cartesian dualism: How could anything physical (or material) be *rational*? How could rationality be mechanized? The fundamental sense of rationality at issue here is the parallelism between the causal and the semantic properties of an individual's thoughts, which can be illustrated by a simple schematic reasoning in accordance with *modus ponens*: John believes *q because* he believes that if *p* then *q* and because he believes *p*. For a pattern of thoughts to be rational in the relevant sense, the causal relations among thoughts must mirror the semantic relations among their contents. Notice that much of the conclusion of Fodor's (1975: 198–202) book, *The Language of Thought*, expresses reasons for skepticism about the scope of the computational approach to mental processes. As Fodor notices, many psychological events may fall outside the scope of the computational approach either because they lack a proper psychological cause or because they do not stand to their psychological

cause in an appropriate computational relation but instead in, for example, some association relation.

The Fodorian trichotomy

As is well known, Chomsky (1975) introduced the notion of modularity in cognitive science as part of his argument that what enables human children to acquire knowledge of the grammar of their native language from primary linguistic data is universal grammar, that is, a modular learning mechanism specific to language learning (and not part of general intelligence). The main issue addressed by the Fodorian notion of modularity is different: Fodor's (1983) main goal is to offer a *nonbehaviorist* alternative to what Fodor calls the "Handsome Cognitivist" denial of the distinction between perception and problem solving (or higher cognition). In a nutshell, the main purpose of Fodorian modularity is to account for the differences between reflexes, perception, and thought or belief-fixation. Unlike thinking, perception is modular. Unlike reflexes, it is computational (and hence representational).

As Sperber (1994) has insightfully noticed, *Modularity of Mind* is a paradoxical title for Fodor's (1983) book since on Fodor's view, modularity is a property only of the *periphery* of the mind (i.e., the input and output systems of the mind) not its core, which, in Fodor's own terms, is Quinean (i.e., holistic) and isotropic (i.e., nonmodular). Arguably, the single most important feature of modularity in Fodor's (1983) sense is informational encapsulation, that is, the fact that the output of a modular information processing system is insensitive to much of the information available to the system.[2] Perceiving (seeing, hearing, touching, or smelling) something is not believing. For example, when seeing a display of the Müller-Lyer visual illusion, one sees two equal line segments as unequal, and the fact that one knows (and therefore believes) that they are in fact equal does not suppress the illusory visual experience. This shows that visual experience is informationally encapsulated.

What must be further noticed is that Fodor's (1983) conception of modular input systems is in fact part of a *tripartite* distinction between central systems, modular input systems, and bodily transducers. (This tripartite distinction has been suggestively, if derogatorily, called "the sandwich model" of the mind by Hurley, 1998, 2008.) Transducers are located on the bodily surface and receive afferent information from the world: They convert information (landing on e.g., the retina) from one medium (e.g., photons) into another medium (e.g., electrons) and relay the information to input systems, which lie at the periphery of the mind. In a nutshell, the core of the mind is constituted by nonmodular central systems, which receive information from modular input systems, which filter information from bodily transducers, which in turn relay afferent information from the world. The same chain of information processing – but in the reverse efferent order – turns an agent's prior intention (stored in her central thought processes) into a motor intention (in a format readable by the mind's modular output systems),

which is then converted by motor transducers at the bodily surface in a format suitable for the execution of bodily movements.

Neurocentrism

The final assumption constitutive of Cartesian materialism is the physicalist thesis that an individual's mind is identical to, or supervenes upon, her brain alone – an assumption that is, of course, inconsistent with Cartesian ontological dualism. In summary, on the Cartesian materialist picture, an individual's mind is protected from both its nonneural bodily environment and its nonbodily environment. The core of the mind – which Fodor (1983) calls central thought processes – is surrounded by a peripheral shell of modular input systems, both of which are securely anchored in the individual's brain (or central nervous system) and protected from the world by a further shell of bodily transducers. Unlike advocates of moderate embodiment, and like advocates of the extended mind thesis, advocates of radical embodiment are motivated by a deep aversion towards a view of the mind secluded from the nonneural bodily environment of the brain. But, unlike advocates of the extended mind thesis, they are willing to endorse the boundary between an agent's embodied mind and its nonbodily environment.

Moderate embodiment

Advocates of moderate embodiment hold that some higher cognitive processes are embodied in the sense that they involve an individual's mental *representations* of some of her body and bodily parts. Unlike advocates of radical embodiment who make possession of an agent's body a constitutive part of her mind, advocates of moderate embodiment stress the contribution of an agent's mental representations of her body to higher cognitive processes. For instance, counting integers smaller than 10 may involve the mental representations of one's own pair of hands and fingers and the motor control and monitoring of hand and finger movements. Studies showing activations of an agent's somatosensory cortex and/or motor cortex during tool use can also be taken as instances of moderate embodiment (cf. Maravita and Iriki, 2004 for a review). Now, as the case of phantom limbs painfully shows, an agent's mental representations of one her amputated limb may be a *misrepresentation* of the presence of a missing bodily part. In this case, the (continued) possession of an agent's bodily part cannot be a necessary and constitutive part of her mind.

One good theoretical example of moderate embodiment is the doctrine that Barsalou (2008) calls *grounded cognition*, whose major goal is to reject the core assumption of concept-*rationalism* that concepts are *amodal* symbols in the language of thought (cf. Fodor, 1975). Grounded cognition endorses instead a two-tiered version of concept-*empiricism*, according to which concepts are (visual, auditory, olfactory, tactile, or motor) *images* and higher cognitive processes are processes of *simulation* (or re-enactment) of basic perceptual and motor processes (cf. Barsalou, 1999, 2008, Gallese and Lakoff, 2005, and Prinz, 2002, 2005. For

critical discussion, see Machery, 2006, 2007, Mahon and Caramazza, 2008, and Jacob, 2012). Thus, while concept-empiricism endorses moderate embodiment, it is entirely consistent with neurocentric assumptions.

One major challenge for moderate embodiment is to spell out the conditions under which a mental representation of an individual's body or bodily parts counts as an embodied representation. Clearly, not any representation of an individual's body (or bodily parts) can count as embodied. For example, a scientific textbook of human anatomy and/or human physiology contains abstract theoretical propositions *about*, and iconic illustrations *of*, human bodies and bodily parts. However, neither abstract theoretical propositions about, nor iconic illustrations of, human bodies should presumably count as embodied representations.

Goldman (2012) and Goldman and de Vignemont (2009) have recently tried to meet this challenge in the domain of social cognition by arguing that what makes a representation of another's action embodied is its bodily *format* (or *code*). On their view, a mental representation derives its bodily format from two joint features: First, it represents the agent's *own* bodily parts from a *first-person* interoceptive (e.g., proprioceptive) perspective (only accessible to the self). Second, the representation is co-opted (exapted or redeployed) for a different purpose (e.g., for representing *another*'s bodily part). Their proposal applies paradigmatically to mirror neuron activity in an observer's brain: Mirroring an agent's goal-directed action (in an observer's brain) is taken to involve a representation of another's bodily action in a *bodily format* because it is the output of a process of redeployment (or exaptation) of a brain mechanism whose original function is to underlie the execution of action in the observer's brain. This process is called "embodied simulation" by Gallese and Sinigaglia, 2011.[3]

While moderate embodiment is the claim that higher cognitive processes may depend on the representations of human bodily parts, radical embodiment is the claim that what makes an individual's cognitive process embodied is the contribution of the *human bodily anatomy* itself to the execution of higher cognitive tasks, namely, of the individual's body to the cognitive process. Thus, moderate embodiment is compatible with the neurocentric assumption that an individual's mind is identical to, or supervenes on, the individual's brain. But radical embodiment is not. To subscribe to radical embodiment is to deny that an individual's mind reduces to the individual's brain and to embrace the claim that an individual's nonneural body is a *constitutive* part of her *mind*. The gap between moderate and radical embodiment is clearly exemplified by de Bruin and Gallagher's (2012) complaint (on behalf of radical embodiment) that "bodily formatted representations are nothing other than brain processes" (99).

From crude to sophisticated radical embodiment

In short, moderate embodiment stresses the role of mental representations of bodily parts in higher cognitive processes. This is why moderate embodiment is attractive to advocates of concept-empiricism: moderate embodiment offers an alternative

to concept-rationalism, in accordance with neurocentric assumptions. By contrast, what is common to all versions of radical embodiment is that an agent's possession of her bodily anatomy is taken to be a constitutive part of her mind, in violation of neurocentric assumptions.

Intelligence without representations

Radical embodiment can be crude or sophisticated.[4] Advocates of crude radical embodiment reject primarily the computational approach to cognition. In Shapiro's (2007: 338) words, proponents of crude radical embodiment advocate "an approach to cognition that departs from traditional cognitive science in its reluctance to conceive of cognition as computational and in its emphasis on the significance of an organism's body in how and what the organism thinks." Advocates of crude radical embodiment are dissatisfied with the computational approach to mental processes for at least two reasons. On the one hand, Shapiro (2004) rejects the computational approach on the grounds that it entails the notorious thesis of multiple realizability, which he further decomposes into two subtheses, which he labels respectively *body neutrality* and the *separability thesis* (i.e., the separability between the computer program and the body that implements the program), both of which are inconsistent with radical embodiment (cf. Shapiro, 2011).

On the other hand, much of the appeal of crude radical embodiment rests on its commitment towards minimizing the load of internal processing and the role of mental representations in the execution of cognitive tasks, as epitomized by the title of Brooks' (1991) famous paper, "Intelligence without representation." Many advocates of crude radical embodiment tend to reject the computational approach to the mind precisely because it entails (or presupposes) a heavy commitment to a representationalist framework. Thus, in accordance with Brooks' (1991) parsimonious recommendation that we should "use the world as its own best model,"[5] Noë (2004) argues that "there is no need to re-present the world on one's own internal memory drive. Off-loading internal processing onto the world simplifies our cognitive lives and makes good engineering sense." In short, mental representations are taken to be dispensable from cognitive scientific explanations. The main alternative to the computational approach to the mind offered by advocates of crude radical embodiment is the direct link between the agent's possession of her body and her ability to perform actions, which advocates of the so-called *enactivist* perspective take in turn to underlie the execution of all cognitive tasks (cf. O'Regan and Noë, 2001 and Gallagher, 2005. See Jacob, 2008 and de Vignemont, 2011, for discussion).

The main challenge for the commitment to action by advocates of crude radical embodiment is that there is more to action than the execution of bodily movements. As philosophers of action have stressed for a long time, an agent's bodily movements count as an action only if they are appropriately caused by the agent's intentions. For example, an agent can execute one and the same hand gesture to frighten a fly or to waive bye-bye to a departing host. Only by representing the

agent's two distinct intentions can one action be distinguished from the other. Furthermore, there is much empirical evidence showing that an agent's motor system is activated in at least two situations in which she fails to perform any overt bodily movement. In such situations, the agent's motor system is activated off-line (see Jeannerod, 2006). On the one hand, some areas of an observer's motor and premotor systems (e.g., mirror neurons) are active when he or she perceives an action performed by another agent. On the other hand, parts of an agent's motor system are active in tasks of motor imagery whereby the agent plans and/or imagines an action, which, for some reason or another, she fails to execute. In fact, humans mentally represent and even plan many actions that they never carry out (cf. Jacob, 2014a and Jacob, in press).

Intelligence without transduction

While the main concern of advocates of crude radical embodiment is to avoid the costs incurred by the postulation of internal mental representations, the main concern of advocates of sophisticated radical embodiment is to undermine the set of boundaries postulated by the Fodorian trichotomy. One way to reject the Fodorian trichotomy is to reject the distinction between central thought processes and modular input and output systems, and to endorse instead the thesis of *massive modularity* (Sperber, 1994, 2002, 2005). In response to the massive modularity thesis, Fodor (2000: 99) has limited the scope of the computational approach to modular input systems.[6] He argues that genuine computational processes are subject to a *locality* constraint: They are local processes defined over the syntactic properties of mental representations. But nondeductive (i.e., abductive and inductive) reasoning processes, which are typical of central thought processes, cannot be subject to the locality constraint because they depend on global contextual nonsyntactic factors. (For a response to Fodor's move, see Sperber, 2002.)

Far from endorsing massive modularity, advocates of sophisticated radical embodiment are skeptical, not only of the application of the modularity thesis to input and output systems, but also of the very notion of *transduction*, both of which stand at the basis of the Fodorian trichotomy. One of the earliest cornerstones of sophisticated radical embodiment is John Haugeland's (1995) paper "Mind embodied and embedded," whose main aim is to articulate a picture of *intelligence without transduction*. Haugeland's critique of the Fodorian trichotomy proceeds in three main steps.

First, Haugeland (1995) argues that the principles of modular design can apply only if and when there exist physical boundaries ("corporeal interfaces") between self-contained components of an information-processing system. For example, the principles of modular design apply to the behavior of a TV set, in which a resistor is a genuinely independent electronic component, such that nothing that happens outside can affect anything that happens inside, and vice versa. Thus, the connecting wires of a resistor constitute a well-defined physical interface with the rest of the system and a resistor is replaceable by a functional equivalent.

Second, in anticipation of arguments for the extended mind thesis (see section 4), Haugeland considers Herbert Simon's (1969: 63–64) well-known example of an ant making its "laborious way across a wind-and-wave-molded beach." Simon's point was that the complexity of the ant's path "is really a complexity in the surface of the beach, not a complexity in the ant." Haugeland (1995: 215–217) argues that, while the structure of the beach is irrelevant to understanding the ant's respiration or immune system, the surface of the beach is an intimate part of the ant's path on the sand. As far as the ant's path in the sand is concerned, the ant and the surface of the beach form a coupled dynamical system. If no relevant physical boundary (no corporeal interface) between the surface of the beach and the ant's body (let alone its nervous system) can be drawn, then for the purpose of explaining the ant's path on the sand, the principles of modular design cannot apply. (But, see Gallistel (1998) for a defense of a computational approach to insect navigation, consistent with modularity assumptions.)

Haugeland extends the critique of modular design from the ant's behavior to human navigation. Not only does a human driver form a tightly coupled dynamical system with her car, but the human driver located in her car also forms a tightly coupled dynamical system with the road on which she is driving: "the internal guidance system and the road itself must be closely coupled" (Haugeland, 1995: 234). Haugeland's point is that "the road itself should be considered an integral part" of the driver's ability to navigate. His claim is not that all, but some, of the structure underlying human intelligent navigation is encoded in the road itself. Even so, there is no relevant "corporeal interface" (or physical boundary) between the road and the driver's cognitive capacities involved in navigation. It follows that the principles of modular design fail to apply to a human driver's ability to flexibly navigate on humanly built roads. Arguably, however, Haugeland thereby considerably underestimates the extent to which the distinctive flexibility exhibited by human drivers depends jointly on the intentionality of the human engineers who designed the road and on the driver's internal ability both to extract spatial information from landmarks posted on the road and to convert spatial information coded in an allocentric frame of reference into information coded in an egocentric frame. In particular, the visual system of humans enables them to write and read (cf. the discussion of the extended mind thesis in the next section.

The last step of Haugeland's critique of the Fodorian trichotomy is meant to undermine the very notion of *transduction*. According to the Fodorian trichotomy (or the *sandwich* model of the mind in Hurley's sense), afferent information from the world cannot affect an individual's brain unless it is converted by bodily transducers into a format that is readable by modular input systems. Nor can an individual's brain cause changes in the world unless modular output systems send efferent information (i.e., motor instructions), which can be decoded by bodily transducers and converted into movement executable by bodily effectors. While Haugeland argues that modularity only applies to the components of an information-processing system if and when they are separated from one another by physical boundaries, he further argues that the very notion of transduction is

entirely committed to the Cartesian separation of an individual's mind from both her body and the world. Bodily transducers are, from Haugeland's standpoint, nothing but a bodily shell around the mind, whose main Cartesian function is to protect the mind from worldly intrusions: Transduction seems tailor made for Descartes' pineal gland. Furthermore, "the idea that there are [motor] instructions is morally equivalent to the idea that there are transducers" (Haugeland, 1995: 223–224). As he puts it at the end of his manifesto, "*if we are to understand mind as the locus of intelligence*, we cannot follow Descartes in regarding it as separable in principle from the body and the world. . . . Mind, therefore, is not incidentally but *intimately embodied and intimately embedded in its world*" (236–237).

On the face of it, the negative content of Haugeland's anti-Cartesian *intuition* is clear enough: The postulation of *boundaries* between an individual's mind, her body, and the world cannot help promote, but only interfere with, our understanding of human intelligence. But, two fundamental questions arise, the first of which is what is the sophisticated radical embodied alternative to the postulation of Cartesian boundaries? The second correlative fundamental question is whether understanding human *intelligence* is a suitable topic of cognitive scientific research. I start with the first question.

It is, I think, fair to say that so far advocates of sophisticated radical embodiment have only scratched the surface of a potentially non-Cartesian conception of human intelligence freed from both modularity and transduction. For example, Haugeland (1995) has tried to sketch an explicit non-Cartesian alternative to the Fodorian trichotomous conception according to which the brain can only cause changes in the world by means of motor instructions that are decoded and converted by bodily transducers into movements executable by bodily effectors. Haugeland's sketch focuses on the notion of *skillful activity* instantiated by human acts of typing, dancing, driving, or speaking. Thus, Haugeland argues that typing a particular letter (e.g., capital "A") with one's fingers does not involve some efferent transduction from the agent's brain to the production of finger movements: It involves instead a *muscular gestalt*. As he puts it (229), "the meaningful (mental) extends all the way to the fingertips . . . and then interfaces to the physical world."

Haugeland's assumption here is that the Cartesian picture of efferent transduction cannot accommodate the bodily complexity and intrinsic interconnectedness exhibited by human skillful activities. For example, an agent's skillful typing of capital "A" on a particular typewriter depends to a large extent on such factors as the particular length of her fingers, the peculiar strength and quickness of her muscles, the shapes of her joints. If so, then it is unlikely that we could ever sort out the respective contribution of each of these factors. Nor could the motor command issued from one agent's brain ever cause another agent, whose fingers, muscles, and joints had a different size, intensity and shape, to skillfully type a token of capital "A" on the same type writer. In a nutshell, Haugeland's non-Cartesian picture of intelligence without transduction seems predicated on his rejection of a *functionalist* construal of the role of bodily movements in skilled human actions.

A related potential alternative to the Fodorian trichotomy involves the notion of *smooth coping* (borrowed from Heidegger's philosophy). As Michael Wheeler (2005: 135) has put it,

> [E]veryday cognition is fundamentally a matter of *smooth coping*. And smooth coping is characterized by circumspection, a form of awareness in which there is no phenomenological distinction to be drawn between subject and object, there is only the experience of the ongoing behavior.

Whichever criterion one chooses for the achievement of smooth coping, it is unlikely to be the human default response to every cognitive challenge. Arguably, humans smoothly cope when they deal under time pressure with *face-recognition*, but not with *multiplications of large integers*. Furthermore, if and when it does occur, "smooth coping" must denote the *output of some underlying cognitive process*, not the process itself. So the question is: what is the psychological process that underlies smooth coping?

While smooth coping is arguably an important part of what Haugeland thinks of as non-Cartesian human intelligence without transduction, that is, purged of any boundaries between an agent's mind, her body and the world, it is not clear to what extent smooth coping is a suitable topic for cognitive scientific research. Furthermore, while a non-Cartesian view of the contribution of both the human body and the nonbodily environment to human intelligence and/or human intelligent behavior might be a topic for philosophical controversy, it is far from clear to what extent, unlike the human language faculty or human vision, human intelligence and human intelligent behavior are suitable topics for cognitive scientific research.[7]

Challenges for the body-centric rejection of neurocentrism

The third fundamental tenet of Cartesian materialism is its commitment to the neurocentric assumption that an individual's mind is identical to, or supervenes on, her brain alone, or that mental processes are brain processes. As I mentioned earlier, unlike moderate embodiment, radical embodiment is committed to rejecting neurocentrism (sometimes called "brain chauvinism"). At least, two questions arise, the first of which is: can one reject neurocentrism without embracing Cartesian ontological dualism? While the extended mind thesis paves the way for a positive answer to this question, the further question arises whether radical embodiment can distinguish itself from the extended mind thesis. I turn to the first question first.

The extended mind thesis

Can one both reject neurocentrism and ontological dualism? A straightforward positive answer to this question is Clark and Chalmers' (1998) *extended mind* thesis,

according to which an individual's mind should not be limited to her brain alone but should instead be construed as a *three*-place relation between her brain, its proximal nonneural bodily environment and its distal nonbodily environment (cf. also Clark, 2008b). The major argument advanced by advocates of the extended mind thesis for the rejection of neurocentrism is what, following Dennett (1996: 177–178), I will call the *off-loading* argument: "the primary source (of our greater intelligence than our nearest relatives) is our habit of off-loading as much as possible of our cognitive tasks into the environment itself – extruding our minds." (Incidentally, it is unclear how one can measure degrees of intelligence across different species.) Dennett (1996: 138) further made the point that older people with memory deficits load their home environments with landmarks that help them solve various everyday tasks (e.g., where to find their keys). As he puts it, "taking them out of their homes is literally separating them from large parts of their minds – potentially just as devastating a development as undergoing brain surgery." Similarly, Clark and Chalmers (1998) have urged that a person's notebook is an integral part of her mind and Chalmers (2008) has claimed that his recently acquired iPhone has become part of his mind. In a nutshell, the point of the offloading argument is that much of the burden of an individual's biological (episodic or semantic) memory in the execution of higher cognitive tasks can be alleviated by storing information in, and retrieving it from, external memory devices, not located within the boundaries of the individual's skull (cf. Jacob, 2012).

While the extended mind thesis is perfectly consistent with the physicalist rejection of Cartesian ontological dualism, the appeal to the off-loading argument raises at least two basic problems, the first of which is that it is not entirely clear to what extent it is sufficient to support the extended mind thesis. To see why, I will make a liberal or extended (no pun intended) use of the notion of a *tool*, which is, I think, very much in line with the current cognitive neuroscientific approach to human action. According to Beck's (1980) widely accepted view of what it takes for an animal to be a tool user, a tool is some unattached environmental object used to effect a change (in the form, position or condition) of another object, another organism, or the user itself. On this view, an animal could not be a tool user unless it had the bodily and motoric resources necessary for retrieving, extracting, tuning, carrying around, monitoring, and possibly storing relevant objects. While there is a clear intuitive distinction between unattached objects which are parts of an animal's distal environment and the animal's attached bodily parts, some selected parts of the animal's own body are clearly necessary for its use of a tool (i.e., an unattached object). For example, birds use their beaks and primates use their hands. Healthy humans use specific bodily parts to perform actions with tools, including cognitive tools such as a compass or an abacus. Furthermore, in line with the basic model of mirror neuron activity, which is taken to code primarily the agent's *goal* (not the agent's bodily movements), we can construe the agent's bodily part as a bodily *tool* selected by the agent's motor system as a means to achieve her goal. Thus, Gazzola et al. (2007) argue that their findings show that "aplasics born without hands mirror the goal of hand actions with their feet."

They thereby seem to assume that whereas healthy agents use their hands, aplasics born without hands use their feet as bodily tools, to achieve the same goals. An agent's bodily limbs turn out to be tools recruited by the agent's motor system.

From the fact that some bodily or nonbodily tool enhances an individual's execution of a cognitive task, it does not immediately follow that the tool itself is *part* of the individual's mind. Presumably, just as Chalmers bought his iPhone before using it, he could intentionally sell it or lend it to someone else, or have it nonintentionally stolen. Now, it does not really make much sense to trade (buy, sell or lend) a part of one's mind. Nor to have it stolen. What the extended mind thesis requires is not merely that some tool play a definite *causal* role in some cognitive process, but that it be a *constitutive* part of the individual's mind. But this is something stronger than what can be established by the off-loading argument alone. To see this, suppose that drinking a cup of coffee or taking a walk on a trail enhances my execution of an arithmetical addition. Does that make either the cup of coffee or the trail a part of my mind? If not, why not?

Secondly, advocates of the extended mind thesis usually fail to recognize that the offloading argument rests on a fundamental *trade-off* between internal brain resources and the use of external memory devices. An individual could simply not off-load the burden of her biological memory onto her nonbodily environment (e.g., into a notebook or an iPhone) unless she had a *literate* brain and knew how to read and write. In literate humans, a small part of their left fusiform gyrus (a small area in the human visual cortex), known as the Visual Word Form Area (VWFA), which underlies the visual perception of the shapes of words, has become a specialized interface between human vision and the human language faculty (cf. Dehaene and Cohen, 2011).

Thus, the leading intuition underlying the extended mind thesis is that storing information in some external memory device may off-load some of the burden of internal cognitive processes (e.g., memory). Now, Clark and Chalmers (1998) present the extended mind thesis as a version of an *externalist* view of the mind, purged of some of the weaknesses of content externalism based on so-called "Twin-Earth" arguments (earlier articulated by Putnam, 1974, and Burge, 1979). In fact, the basic off-loading argument for the extended mind thesis significantly differs from standard arguments for content externalism based on Twin-Earth thought experiments, at least in one crucial interesting respect.

Content externalism is the view that the contents of an individual's psychological states (e.g., her beliefs) depend not merely on intrinsic properties of the individual's brain, but on properties instantiated in the individual's environment as well. Standard Twin-Earth arguments for content externalism presuppose that an individual's brain and/or cognitive resources are kept constant throughout variations in the individual's environment relevant to establishing the contribution of the social or nonsocial environment to the individuation of the content of the individual's psychological state. Clark and Chalmers (1998: 9) have pointed out that content externalism is a passive version of externalism because the relevant external features are distal and historical, "playing no direct role in driving the

cognitive process in the here-and-now." By contrast, the extended mind would constitute an *active* version of externalism.[8]

But the off-loading argument for the extended mind thesis rests on a trade-off, which casts some doubt as to whether it is a genuine version of externalism: an individual could not store some information in, and retrieve it from, some external memory device (e.g., a notebook), unless her brain enabled her to read and write. If so, then the off-loading argument for the extended mind thesis could not go through unless a fundamental change in the individual's brain resources was posited before and after the individual becomes able to alleviate the burden of her biological memory (for further discussion cf. Jacob, 2014b). If so, then a book, a notebook or a computer (located outside the individual's brain and skin) could only be part of the individual's *extended* mind if the individual's brain contained the *internal* cognitive resources necessary for storing and extracting information from the book, the notebook or the computer. In a nutshell, the individual's brain must be a *literate* brain.

Not only does recognition of this trade-off call into question the extent to which the extended mind thesis supports a genuine version of externalism, but it further gives rise to what Clark (2002: 70) has called the *paradox of active stupidity*, namely "the idea that making the moves that sculpt the environment so as to allow *cheap* problem solving itself requires *expensive*, advanced, design-oriented cogitation. The nasty upshot being that only *clever* brains could make their worlds smart so that they could be dumb in peace." It is, in my opinion, far from clear how the extended mind thesis can meet this pair of challenges. I now turn to the second question, namely, the question whether advocates of radical embodiment can reject neurocentrism without endorsing the extended mind thesis.

The body-centric rejection of neurocentrism

Advocates of radical embodiment do not mean to argue that nonbodily tools (e.g., an iPhone) are parts of an agent's mind. They mean to argue instead that only an agent's nonneural bodily anatomy is part of her mind. To the extent that the basic argument against neurocentrism is the off-loading argument (used by advocates of the extended mind thesis), advocates of radical embodiment face a pair of hard challenges. First, radical embodiment must be controversial while keeping its distance from moderate embodiment. It must keep its distance from moderate embodiment because moderate embodiment is compatible with neurocentrism. So, the first problem for radical embodiment is to reject neurocentrism without collapsing into what I called in the introduction noncontroversial embodiment.

The first challenge faced by advocates of radical embodiment is that there is a slippery slope between controversial and noncontroversial versions of embodiment. As I said, it is uncontroversial that an agent's execution of some cognitive task (e.g., reasoning or computing a multiplication) causally depends on the possession of a number of physiological and bodily systems and/or organs such as the

digestive, respiratory and cardiovascular systems. It is easy to slide from claims about something being a causally enhancing condition for something else to a claim about identity. But it is clearly a mistake. To see why, consider the controversial identity claim embraced by some philosophers that mind is life or that "where there is life there is mind" (Thompson, 2007, 2011). For example, one of Noë's (2009) primary concerns is to reject the neurocentric (or brain chauvinistic) assumption that an individual's mind is identical to, or supervenes on, her brain alone. On Noë's (2009: 42) view, having a brain is necessary but not sufficient for having a mind: "only creatures with the right kinds of brains can have certain kinds of experiences, and to events in consciousness there doubtless correspond neural events. But there are external correlates of consciousness too." However, as a result of his further acceptance of the identity claim that mind is life, Noë is led astray to granting a mind to a bacterium: "The mind of the bacterium does not consist in something about the way it is internally organized. . . . The mind of the bacterium, such as it is, consists in its form of engagement with and gearing into the world around it. Its mind is its life" (42). But clearly, one cannot both hold that a bacterium has a mind and that only creatures with the right kinds of brain can have psychological experiences. What has potentially gone wrong here is that it is one thing to grant that life is a necessary condition for having a mind.[9] Another thing is to identify mind and life.

So the first challenge for advocates of radical embodiment is to argue that an agent's nonneural bodily parts are not mere bodily tools at the service of the agent's brain, but that they are parts of the agent's mind. For example, the task for advocates of radical embodiment is to show that if an agent uses, for example, the five fingers of her right hand in a task of counting integers smaller than five, then her right hand with five fingers is not a bodily tool that plays a causally efficacious role in her execution of the arithmetical task, but a genuine constitutive part of her mind. If the five fingers of an agent's right hand are not just causal enabling conditions for cognitive processing, but constitutive parts of her mind, then why should not her digestive system, her sweat or sudoriferous gland system, and her cardiovascular system be too? Conversely, if the latter are not, then why should an agent's right fingers be?

Secondly, advocates of the extended mind thesis construe an individual's mind as a *three*-place relation between her brain, its proximal nonneural bodily environment and its nonbodily environment. But advocates of radical embodiment endorse what Clark (2008a) calls a *body-centric* view of an individual's mind, which they construe as a *two*-place relation between the individual's brain and its proximal nonneural bodily environment, not as a three-place relation (involving also the nonbodily environment of the brain). The basic challenge for radical embodiment is to justify the restrictive body-centric scope of its rejection of neurocentrism. Why limit the nonneural constituents of an agent's mind to the bodily environment of her brain (what Clark, 2008a calls "the flesh"), at the expense of its nonbodily environment? Why should an agent's mind include bodily tools, at the expense of nonbodily tools?

One possible line of defense of body-centrism is to appeal to some unique features of the sense of *ownership* that applies selectively to an agent's *bodily parts*, at the expense of an agent's nonbodily tools, some of which are explored in depth by de Vignemont (in preparation). (It is an open question whether the sense of bodily ownership is restricted to skeletal bodily parts or also applies to nonskeletal bodily parts, e.g., guts.) The first feature of an agent's sense of bodily ownership is that it seems independent from her sense of agency in a unique way. Arguably, whether innate or not, an agent's sense of bodily ownership does not depend on her executing any bodily action: I feel my hand as my own even when I do not move it, but you move it (cf. Tsakiris et al., 2010). By contrast, I could not own some nonbodily tool unless I had previously performed some action that enabled me to acquire it by, for example, taking possession of it first.[10] Furthermore, a human agent can trade any nonbodily tool that she owns, relinquish her ownership and transfer it to some other agent. But she cannot, without undergoing surgery, rid herself of some of her bodily part.

Secondly, as is well known, unlike the transfer of some nonbodily tool to someone else, the amputation of a limb gives rise to the rich and painful phenomenology of phantom limbs. As de Vignemont (chapter 8, this volume) points out, it seems intuitively clear that we *care* about our bodily limbs in special ways: We frequently use nonbodily tools in ways that we would never use our limbs, for example, in mending a fire. Even chimpanzees have been reported by Povinelli et al. (2010) to use a tool, not their hands, to remove the cover of a box when they perceived the object in the box as potentially hazardous, and with their hands, when they perceived that the box contained food.

However, the experimental investigation of the effects of tool-use in both nonhuman primates and humans raises a significant challenge for a strong functional demarcation between bodily ownership and the use of nonbodily tools. For example, for many years, Iriki and colleagues have trained macaque monkeys to use a rake in order to retrieve objects in extra-personal space too far to be reached with their unaided arm and hand. They recorded bimodal (visual and somatosensory) neurons in the intra-parietal cortex before and after training. They found an expansion of the neurons' receptive fields, which they interpret as evidence that the rake has been incorporated into the monkeys' body schema, as if the monkeys' own effectors were elongated to the tip of the tool (cf. Maravita and Iriki, 2004 for review). More recently, Cardinali et al. (2009) have reported that after using a tool, the kinematics of humans' arm movements is significantly modified. They argue that the observed kinematic changes in arm movements induced by tool-use can be taken to reflect somatosensory changes in the agent's representation of her own arm, that is, body schema.

Such findings are entirely consistent with neurocentrism and can be easily accommodated by moderate embodiment. They are also consistent with a functionalist approach to the contribution of both nonbodily and bodily tools to the execution of cognitive tasks, in accordance with the extended mind thesis. They can be taken to support the extended functionalist picture of the body which is

just, according to Clark (2008a: 56–57), "one element in a kind of equal-partners dance between brain, body and world, with the nature of the mind fixed by the overall balance thus achieved." But they raise a challenge for the strong dichotomy between bodily and nonbodily tools, which lies at the core of radical embodiment.

Concluding remarks

The main goal of this chapter has been to assess the costs incurred by advocates of radical embodiment whose own main agenda is to deny the need for boundaries between an agent's mind and her nonneural bodily organs in cognitive scientific research. Much of this chapter rests on a pair of simple distinctions, both of which are often overlooked by advocates of embodied cognition. One is the distinction between uncontroversial and controversial claims about the role of nonneural bodily organs and systems in the execution of higher cognitive functions. The other is the distinction between two ways an agent's execution of a cognitive process may be said to be embodied: by stressing the role of either the agent's possession of her nonneural bodily anatomy (radical embodiment) or of the agent's ability to mentally represent some of her relevant bodily parts (moderate embodiment). I have further subdivided radical embodiment into cruder (e.g., enactivist) versions that seek to eliminate mental representations from the tool-kit of cognitive scientific research and more sophisticated versions that seek alternatives to the Fodorian trichotomy between central processes, modular input systems and bodily transducers. So far as I can see, neither version of radical embodiment has offered a coherent alternative. Furthermore, radical embodiment is crucially committed to rejecting the neurocentric assumption that an individual's mind is identical to, or supervenes on, her brain alone. So far as I can tell, the main current argument against neurocentrism is what I called the *off-loading* argument offered by advocates of the extended mind thesis, according to which nonbodily tools and bodily tools are constitutive parts of an agent's mind. But while the off-loading argument alone is not sufficient to support the extended mind thesis (or so I have argued), advocates of radical embodiment face the further challenge of explaining why, unlike advocates of the extended mind, they are not willing to allow nonbodily tools as constitutive parts of an individual's mind.[11]

Notes

1 Representationalism in the sense relevant to this chapter (as entailed by the computational approach to mental processes) is related but different and weaker than representationalism (or intentionalism) in the sense advocated by philosophers of mind in Brentano's tradition who argue that the phenomenal properties (or phenomenology) of sensory experiences are reducible to their representational properties.

2 In addition to informational encapsulation, Fodor (1983) lists eight further features of modularity: domain specificity, the mandatory character of modular processing, the limited conscious access to the output of modular processes, the speed of input systems, the shallow output of modular processes, the fixed neural architecture of modular processes, specific patterns of breakdown of modular processes, and specific patterns of ontogenetic development.

3 For further appraisal of embodied simulation in the context of mirror neuron activity, cf. Jacob (2013a–b).
4 As it turns out, advocates of radical embodiment can embrace both the crude and the sophisticated versions.
5 Incidentally, as noticed by Fodor (2009), it is far from clear how a mindless world could be a model of (or represent) itself.
6 As I mentioned earlier, at the end of his (1975) book, Fodor expressed skepticism about the scope of the computational paradigm in psychology. His (2000) response to massive modularity can be seen as reinforcing his earlier skepticism.
7 Much of the inspiration for the attempt by advocates of radical embodiment to reject the Fodorian trichotomy and the boundaries of bodily transducers and to turn to smooth coping as an alternative picture of human intelligence derives from Dreyfus's (1972) critique of classical artificial intelligence (AI).
8 "Because [the relevant external features] are coupled with the human organism, they have a direct impact on the organism and on its behavior . . . Concentrating on this sort of coupling leads us to an *active externalism*, as opposed to the passive externalism of Putnam and Burge" (Clark and Chalmers, 1998: 9).
9 If so, then no nonbiological system should be granted psychological states.
10 Or, unless one of my ancestors had performed some action that enabled her to own it, and I legally inherited it from her.
11 I am grateful to Yann Coello for inviting me to the International Symposium on Vision, Action and Concepts, Behavioral and Neural Bases of Embodied Perception and Cognition, which he organized in Lille (October 28–30, 2013). I am also grateful to comments by Frédérique de Vignemont on my paper and to remarks on my presentations in Lille and in the Philosophy Department at Central European University in Budapest made in particular by Larry Barsalou, Hanoch Ben-Yami, Stephen Butterfill, Kati Farkas, Chris Frith, and Dan Sperber. Finally, I wrote this paper unaware of Hohwy's (2014) recent paper rooted in the basic assumption that the brain is an organ for predicting error and whose views on embodied cognition and the extended mind turn out to be very congenial to mine.

References

Alsmith, A.J.T., and de Vignemont, F. (2012). Embodying the mind and representing the body. *Review of Philosophy and Psychology, 3*, 1, 1–13.

Barsalou, L. W. (1999). Perceptual symbol systems. *The Behavioral and Brain Sciences, 22*, 577–609.

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

Beck, B. (1980). *Animal tool behavior: The use and manufacture of tools*. New York: Garland Press.

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

Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy, 4*, 1, 73–121.

Cardinali, L., Frassinetti, F., Brozzoli, C., Urquizar, C., Roy, A. C., and Farnè, A. (2009). Tool use induces morphological updating of the body schema. *Current Biology, 19*, 12, R478.

Chalmers, D. (2008). Foreword to A. Clark, *Supersizing the mind, embodiment, action, and cognitive extension*. Oxford: Oxford University Press.

Chomsky, N. (1975). *Reflections on language*. New York: Pantheon.

Clark, A. (2002). Minds, brains and tools. In Clapin, H. (Ed.), *Philosophy of mental representations*. Oxford: Oxford University Press.

Clark, A. (2008a). Pressing the flesh: A tension in the study of the embodied, embedded mind? *Philosophy and Phenomenological Research, 76*, 1, 37–59.

Clark, A. (2008b). *Supersizing the mind, embodiment, action, and cognitive extension*. Oxford: Oxford University Press.

Clark, A., and Chalmers, D. (1998). The extended mind thesis. *Analysis, 58*, 1, 7–19.
de Bruin, L., and Gallagher, S. (2012). Embodied simulation, an unproductive explanation: Comment on Gallese and Sinigaglia. *Trends in Cognitive Sciences, 16*, 2, 98–99.
de Vignemont, F. (2011). A mosquito bite against the enactive approach to bodily experiences. *Journal of Philosophy, 108*, 4, 188–204.
de Vignemont, F. (in preparation). *Mind the body.*
Dehaene, S., and Cohen, L. (2011). The unique role of the visual word form area in reading. *Trends in Cognitive Sciences, 15*, 6, 254–262.
Dennett, D. (1991). *Consciousness explained.* New York: Little Brown.
Dennett, D. (1996). *Kinds of mind.* New York: Basic Books.
Dreyfus, H. (1972). *What computers can't do.* New York: Harper & Row.
Fodor, J. A. (1975). *The language of thought.* New York: Crowell.
Fodor, J. A. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.
Fodor, J. A. (1987). *Psychosemantics.* Cambridge, MA: MIT Press.
Fodor, J. A. (1994). *The elm and the expert.* Cambridge, MA: MIT Press.
Fodor, J. A. (1998). *Concepts.* Oxford: Oxford University Press.
Fodor, J. A. (2000). *The mind does not work that way: The scope and limits of computational psychology.* Cambridge, MA: MIT Press.
Fodor, J. A. (2009). Where is my mind? *Times Literary Supplement.*
Gallagher, S. (2005). *How the body shapes the mind.* Oxford: Oxford University Press.
Gallese, V., and Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology, 22*, 3–4, 455–479.
Gallese, V., and Sinigaglia, C. (2011). What is so special about embodied simulation? *Trends in Cognitive Sciences, 15*, 11, 512–519.
Gallistel, C. R. (1998). Brains as symbol processors: The case of insect navigation. In Sternberg, S., and Scarborough, D. (Eds.), *Conceptual and methodological foundations* (Vol. 4, *An invitation to cognitive science*, 2nd ed., pp. 1–51). Cambridge, MA: MIT Press.
Gazzola, V., van der Worp, H., Mulder, T., Wicker, B., Rizzolatti, G., and Keysers, C. (2007). Aplasics born without hands mirror the goal of hand actions with their feet. *Current Biology, 17*, 1235–1240.
Goldman, A. (2012). A moderate approach to embodied cognitive science. *The Review of Philosophy and Psychology, 3*, 1, 71–88.
Goldman, A., and de Vignemont, F. (2009). Is social cognition embodied? *Trends in Cognitive Sciences*, *13,* 4, 154–159.
Hacking, I. (1998). Canguilhem amid the cyborgs. *Economy and Society, 27*(2/3), 202–216.
Haugeland, J. (1995). Mind embodied and embedded. In Houng, Y.-H., and Ho, J.-C. (Eds.), *Mind and cognition* (pp. 3–38). Taipei, Taiwan: Academia Sinica.
Hornsby, J. (1986). Physicalist thinking and conceptions of behaviour. In Pettit, P., and McDowell, J. (Eds.), *Subject, thought and context* (chap. 3, pp. 95–115). Oxford: Oxford University Press.
Horst, S. (2009). The computational theory of mind. In *Stanford encyclopedia of philosophy.* http://plato.stanford.edu/entries/computational-mind/
Hohwy, J. (2014). The self-evidencing brain. *Noûs.* doi:10.1111/nous.12062
Hurley, S. (1998). *Consciousness in.*, Cambridge, MA: Harvard University Press.
Hurley, S. (2008). The shared circuit model (SCM): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behavioral and Brain Sciences, 31*, 1, 58.
Jacob, P. (1997). *What minds can do.* Cambridge: Cambridge University Press.
Jacob, P. (2008). The scope and limits of enactive approaches to visual experiences. *Perception, 47*, 446–461.

Jacob, P. (2012). Embodying the mind by extending it. *The Review of Philosophy and Psychology, 3*, 1, 33–51.

Jacob, P. (2013a). Embodied cognition, communication and the language faculty. In Coello, Y., and Bartolo, A. (Eds.), *Language and action in cognitive neuroscience* (pp. 3–29). London: Psychology Press.

Jacob, P. (2013b). How from action-mirroring to intention-ascription. *Consciousness and Cognition, 22*, 3, 1132–1141.

Jacob, P. (2014a). The legacy of Marc Jeannerod. *Neuropsychologia, 55*, 3–5.

Jacob, P. (2014b). Intentionality. In *Stanford encyclopedia of philosophy.* http://plato.stanford.edu/entries/intentionality/

Jacob, P. (2015). Action-based accounts. In Matthen, M. (Ed.), *The Oxford Handbook in the Philosophy of Perception* (pp. 217–236), Oxford: Oxford University Press.

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

Machery, E. (2006). Two dogmas of neo-empiricism. *Philosophy Compass, 1*, 4, 398–412.

Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition, 104*, 19–46.

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

Maravita, A., and Iriki, A. (2004). Tools for the body (schema). *Trends in Cognitive Sciences, 8*, 2, 79–86.

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

Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

Noë, A. (2009). *Out of our heads*. New York: Hill and Wang.

O'Regan, K., and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*, 5, 939–101.

Povinelli, D. J., Reaux, J. E., and Frey, S. H. (2010). Chimpanzees' context-dependent tool use provides evidence for separable representations of hand and tool even during active use within peripersonal space. *Neuropsychologia, 48*, 243–247.

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

Prinz, J. (2005). The return of concept empiricism. In Cohen, H., and Lefebvre, C. (Eds.) *Categorization and Cognitive Science* (pp. 679–699). Boston: Elsevier.

Prinz, J. (2008). Is consciousness embodied? In Robbins, P., and Aydede, M. (Eds.) *The Cambridge handbook of situated cognition* (pp. 419–437). New York: Cambridge University Press.

Putnam, H. (1974). The meaning of "meaning." In Putnam, H. (Ed.), *philosophical papers, mind, language and reality* (vol. 2, pp. 215–271). Cambridge: Cambridge University Press.

Putnam, H. (1994). Sense, nonsense, and the senses: An inquiry into the powers of the human mind. *The Journal of Philosophy, 91*(9), 445–517.

Rowlands, M. (2010). *The new science of the mind*. Cambridge, MA: MIT Press.

Shapiro, L. (2004). *The mind incarnate*. Cambridge, MA: MIT Press.

Shapiro, L. (2007). The embodied cognition research programme. *Philosophy Compass, 2*, 2, 338–346.

Shapiro L. (2011). *Embodied cognition*. New York: Routledge Press.

Simon, H. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.

Sperber, D. (1994). The modularity of thought and the epidemiology of representations. In Hirschfeld, L., and Gelman, R. (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 39–67). Cambridge: Cambridge University Press.

Sperber, D. (2002). In defense of massive modularity. In Dupoux, E. (Ed.), *Language, brain and cognitive development: Essays in honor of Jacques Mehler* (pp. 47–57). Cambridge, MA: MIT Press.

Sperber, D. (2005). Modularity and relevance: How can a massively modular mind be flexible and context-sensitive? In Carruthers, P., Laurence, S., and Stich, S. (Eds.), *The innate mind: Structure and contents* (pp. 53–68). Cambridge: Cambridge University Press.

Thompson, E. (2007). *Mind in life, biology, phenomenology and the sciences of mind*, Cambridge, MA: Harvard University Press.

Thompson, E. (2011). Précis of *Mind in Life: Biology, Phenomenology, and the Sciences of Mind. Journal of Consciousness Studies, 18*, 5–6, 10–22.

Tsakiris, M., Longo, M. R., and Haggard, P. (2010). Having a body versus moving your body: Neural signatures of agency and body-ownership. *Neuropsychologia, 48*, 2740–2749.

Wheeler, M. (2005). *Reconstructing the cognitive world: The next step.* Cambridge, MA: MIT Press.