THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?

Citation for published version:

Clark, A 2008, 'Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?', *Philosophy and Phenomenological Research*, vol. 76, no. 1, pp. 37-59. <https://doi.org/10.1111/j.1933-1592.2007.00114.x>

Digital Object Identifier (DOI):

[10.1111/j.1933-1592.2007.00114.x](https://doi.org/10.1111/j.1933-1592.2007.00114.x)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Philosophy and Phenomenological Research

Publisher Rights Statement:

© Clark, A. (2008). Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?. *Philosophy and Phenomenological Research*, 76(1), 37-59. 10.1111/j.1933-1592.2007.00114.x

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Pressing the Flesh: Exploring a Tension in the Study of the Embodied, Embedded Mind*.

Andy Clark

Abstract

Much ink has lately been spilled in the name of Embodied Cognition. Mind itself, we are now frequently told, is an intrinsically embodied and environmentally embedded phenomenon. But there is a *prima facie* tension, or so I shall argue, between two strands of thought prominent in this recent literature. One of those strands depicts the fine details of a creature's embodiment as a major constitutive constraint on the nature of its mind: a kind of new-wave body-centrism. The other depicts the body as just 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: a kind of extended functionalism (now with an even broader canvas for multiple realizability than ever before). Where some theorists of embodiment see the contribution of the flesh as both special and at least partially constitutive of certain mental states, some theorists of embedding see the contribution as instrumental, with the body acting as the contingent bridge to new forms of extended functional (computational and representational) organization. The main goal of the present paper is to display this subterranean tension. Along the way, I isolate some specific cases for which a constitutive relation between embodiment and mentality seems most plausible, and begin to scout the space of possible ways of reconciling the two approaches. More positively, but rather speculatively, the paper ends by depicting the body itself as whatever plays a certain role in a (possibly extended) functional organization. This delivers a clear account of why the body matters, but may require some revisions in our common understanding of the space of possible embodied intelligences.

1. Embodiment and Embedding: The Very Ideas

Recent years have seen an explosion of work, both in philosophy and across the many sub-disciplines of Cognitive Science, that is now typically glossed as belonging to the investigation of the mind as 'embodied and environmentally embedded'ⁱ. The phrase 'mind as embodied and embedded' seems to have been coined by John Haugeland in a similarly titled paper that was circulating widely in the early 1990'sⁱⁱ and that later appeared as Haugeland (1998)). There, Haugeland writes that:

“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...Broader approaches, freed of that prejudicial commitment, can look again at perception and action, at skillful involvement with public equipment and social organization, and see not principled separation but all sorts of close coupling and functional unity... Mind, therefore, is not incidentally but *intimately* embodied and *intimately* embedded in its world.”

Haugeland (1998) p. 236–237

Intimacy, however, is a famously slippery beast. What claim or claims concerning the role and importance of the body and the environment lie at the centre of recent work on the embodied, embedded mind? In recent years, some of the many projects developed under this broad banner have included: work on externalism and the nature of psychological explanation (Wilson (2004)); work on 'active externalism' and the extended mind (Clark and Chalmers (1998)); work on 'sensorimotor accounts of perception' (O'Regan and Noë (2001)); work on environment-involving accounts of perception, memory, thought and language (Rowlands (1999)); work on the interdependence of conscious perception and action (Hurley (1998)); work on deictic pointers and active vision (Ballard et al (1997)); and work on the complementarity between biological and technological resources

(Clark (2003)). Given this surface diversity, it seems fair to ask what, if anything, forms the deep theoretical core of the embodied, embedded approach?

It helps, at the outset, to put aside some possible readings that are simply trivial or uninteresting. For example, it is obvious enough that the concept of chair in some way reflects facts about the kinds of shape that happen to afford (typical human) sitting. It is also obvious that much daily thought and reason is informed by what we sense and what we do, and that sensings and doings are paradigmatically embodied acts. Creatures with very different bodies to our own won't (*ceteris paribus*— and more on that later) be prone to think quite the same things in the same circumstances as us. Or take 'knowing how to row'. Knowing how to row, for the average human being, involves knowing how to do certain things with oars and with your arms and legs. What is known is not just how to row, but (if you will) how to row using tools *like this* and a body *like this*. None of this is particularly interesting, either as philosophy or science. The question is whether attention to details of embodiment and embedding contributes something important and previously unexpected to an understanding of mind and reason, not whether it contributes at all.

One way, though probably not the only way, to unpack this idea of embodiment as contributing something 'important and unexpected' is to depict bodily structures or bodily actions as somehow constitutively involved in the determination of aspects of our mental life. Ned Block (personal communication, and see also Block (2001)) suspects that in many cases the actual contribution of the body and world is causal rather than constitutive. He thus accuses many of the research programs gathered under the 'embodied, embedded' banner of committing what he dubs a 'causal/constitutive error'. In what follows, I shall develop two (mutually consistent) responses to this worry. The first is to point out some ways in which 'merely' causal involvement might itself reveal unexpected and interesting things about the mind. The second is to point to an important but restricted range of cases where the strong constitutive claim seems at least *prima facie*

plausible. Given this strategy, it might reasonably be objected that despite being a staple of contemporary philosophical exposition, the causal/constitutive distinction itself is by no means unproblematic. I agree that this is so, but think that the general challenge (which is simply to say more about the nature of the putative role of embodiment and embedding) is nonetheless a fair one. For present purposes, then, I hope we can rest content with a fairly simple gloss according to which the core distinction is between the body (or embodied action) acting as a mere causal factor, channeling and timing inputs and outputs in ways that (of course) impact the flow of experience and mental content, and the more challenging image of the body (or better, embodied activity) as somehow conceptually intertwined with specific mental contents or experiences themselvesⁱⁱⁱ. As we examine specific examples, this contrast should become increasingly clear. One characteristic marker of the very strongest forms of conceptual intertwining is that it should be strictly speaking unimaginable (once the cases are properly presented) that the content or experience remain as it is while certain key details of embodiment or of (perhaps merely potential) embodied action differ. Most importantly of all, however, we will also begin to spot some interesting intermediate territory, in which the details of embodiment, though not conceptually intertwined with mentality in such a strong fashion, play the important role of causally enabling or implementing aspects of the more abstract functional organizations upon which certain mental states and processes supervene.

There is also a large complex of traditionally externalist philosophical theorizing that, while certainly in pursuit of richly constitutive claims, remains orthogonal to the concerns of (most of^{iv}) the theorists of embodiment and embedding. Such philosophical externalism concerns the constitutive role that an embedding environment plays in determining the content of some (or perhaps even all) mental states. But that determination is, in a certain sense, merely semantic. In such cases, the external features are thought to play a role in fixing meanings, but do not appear as part of the local computational and bodily mechanisms

driving environmentally engaged action. By contrast, most theorists of embodiment and embedding focus on ways in which features of the body, of embodied action, or of the environment may play an active^v. information processing role (a clear example is Haugeland's (1998) appeal to the role of the actual road in an agent's 'knowing the way to San Jose').

Having put the trivial and the orthogonal matters aside, what other claim or claims concerning the nature and importance of embodiment and embedding are to be found? Some recent attempts to sort and distinguish various claims and options include M. Wilson's (2002) six-way classification and Shapiro's (2004) three-way division^{vi}. I propose, however, to be rather less delicate than either of these and simply to highlight (what seem to me to be) two rather different central claims that organize and inform the bulk of this recent literature.

The first claim is that the body can be the *instrument* by means of which we create larger functional (computational and representational) wholes, upon which aspects of mind can properly supervene. Call this *Extended Functionalism* (EF).

The second claim is that the specific details of human embodiment, embedding or embodied action make a direct, active and in some sense *constitutive* contribution to at least some of our mental states and properties. Call this the *Constitutive Contribution Claim* (CCC)

The two stories (I hesitate to call them theses at this early stage of philosophical and scientific investigation) are, at least as stated, mutually consistent and even superficially quite similar. The first, EF, claims that, courtesy of our capacities of embodied action, larger systemic wholes, incorporating brains, bodies, the motion of sense organs, and (under some conditions) the information-bearing states of non-biological props and aids, may sometimes constitute the *mechanistic supervenience base* for specific mental states and processes. The second, CCC, claims that specific features of our embodiment are somehow conceptually implicated

the explanation of certain mental states. But despite the superficial similarities (each claims that body and world may play a profound and active role in human cognition) these two stories are both distinct and may, in certain circumstances, come into conflict. Revealing and exploring this hidden tension is the main aim of the present paper.

2. On (Principled) Body–Centrism

Let's take the second story (the one about a constitutive contribution) first. There are various ways to try to make this case, but I'll consider only three. The first appeals to the general role of embodied action in sensing and processing, the second to its role in the determination of certain rather specific forms of conscious experience, and the third to its role in limning the space of thought and reason.

Shapiro (2004)^{vii} pursues the first of these routes as part of an argument against what he dubs *the separability thesis* (ST). According to ST a humanlike mind could perfectly well exist in a very nonhumanlike body. Against ST, Shapiro urges us to embrace what he calls *the embodied mind thesis* (EMT) which holds that “minds profoundly reflect the bodies in which they are contained” (op cit p.167).

Why reject ST? One reason, Shapiro tells us, turns on quite basic facts about sensing and processing. Human vision, for example, involves a great deal of sensor movement. We move our heads to gain information about the relative distances of objects, since nearer objects will (courtesy of parallax effects) appear to move the most. Such movements, Shapiro argues, are not simply an aid to vision. They are part and parcel of the visual processing itself. They are “as much a part of vision as the detection of disparity or the calculation of shape from shading” (op cit p.188). Similar points can be made about audition and the placement of the ears on the head. Rebutting the suggestion that this is merely to state the obvious dependence of perceptual processes on bodily structure, Shapiro comments that:

“The point is deeper—that psychological processes are *incomplete* without the body’s contributions. Vision for human beings is a process that includes features of the human body. [] This means that a description of various perceptual capacities cannot maintain body–neutrality and it also means that an organism with a non–human body will have non–human visual and auditory psychologies” Shapiro (2004) p.190

Body–neutrality, for Shapiro, is the idea that “characteristics of bodies make no difference to the kind of mind one possesses” and is associated with the idea that “mind is a program that can be characterized in abstraction from the kind of body/brain that realizes it” (both quotes op cit p.175). Work on the role of bodily movements in visual processing suggests, according to Shapiro, that body–neutrality fails and that human–style vision requires a human–style body. As I read it, the force of this claim must be at least that 'human–style' vision and 'human–style' embodiment stand in some relation that is more than contingent. Our kind of seeing, understood as a psychological process, requires, or at any rate profoundly involves (see note 11 for some caveats) our kind of body.

Another body (sic) of research that appears to contest claims of body–neutrality, at least regarding the contents of perceptual awareness, is spearheaded by Alva Noë and Kevin O’Regan, and forms part of the development of an ‘enactive’ (see Varela, Thompson and Rosch (1991)) approach to perception. The central claim is that perception is nothing other than implicit knowledge of so–called ‘sensorimotor contingencies’, that is to say, knowledge of the effects of movement on sensory stimulation. The point is not that (trivially) what we see depends on what we do, but that seeing itself is all about our implicit knowledge of the likely effects of our own bodily motions (especially motions of the sensors). Space precludes a full rehearsal of this rich and challenging view, but the conclusions bear repeating:

“If perception is in part constituted by our possession and exercise of bodily skills...then it may also depend on our possession of the sorts of bodies that can encompass those skills, for only a creature with such a body could have those skills. To perceive like us, it follows, you must have a body like ours.” (Noë (2004) p.25)

There is no doubt, in this case, concerning the intended constitutive force of the story. But the waters become muddied when Noë offers an immediate illustration of this claim. As we take off in an airplane, Noë notes, it can look as if the nose lifts up relative to our field of vision. But this is not so, since we are (of course) rising in perfect synchrony with the plane. The nose lifts, but not relative to our visual plane! The illusion occurs because our vestibular system detects the alteration in bodily orientation, and this information impacts our visual experience. According to Noë:

“..the example illustrates the way the character of our visual experience depends on our embodiment, that is, on idiosyncratic aspects of our sensory implementation”
(op cit p.27)

This example, however, fails to do justice to the sensorimotor story. For while the example certainly illustrates an idiosyncratic contribution of sensory implementation, it does not seem, on the face of it, to be tracking any plausibly constitutive connection.

O'Regan and Noë (2001, p.941) offer another, more central, example, this time of a non-illusory nature. It concerns the experience of seeing a long straight horizontal line. This experience involves, according to O'Regan and Noë, knowledge of the peculiar sensorimotor contingencies characteristic of an eye like mine, with a spherically shaped retina, encountering such a line. Thus:

“If you are looking at the midpoint of a horizontal line, the line will trace out a great arc on the inside of your eyeball. If

you now switch your fixation point upwards the curvature of the line will change. Represented on a flattened—our retina the line would now be curved. In general, straight lines on the retina distort dramatically as the eyes move.."

O'Regan and Noë (2001, p.941)

O'Regan and Noë's interesting suggestion is that far from presenting a problem in need of fixing downstream, such signature effects provide the means by which we actually recognize the line as being visually presented and as being straight and horizontal. My experience of seeing the line as horizontal, according to O'Regan and Noë, is thus partially built out of my implicit knowledge of the various contingencies linking specific bodily (or ocular) actions to new sensory inputs.

Another such contingency concerns what happens if we move our eyes along a long horizontal line rather than perpendicular to it. As we move along, the set of photoreceptors stimulated on the eye remains the same (as the line is self-similar under translation). Now, the retinal (and indeed the cortical) representation does not change. This 'law of sensorimotor invariance', unlike the previous one, tracks an intrinsic property of straight lines and would apply regardless of the details of different sensory organs and wiring diagrams (op cit p.942). Even a being with a corrugated or otherwise differently shaped retina would, O'Regan and Noë observe, be subject to this invariance in input signals deriving from sideways visual motion along an endless horizontal line. Conscious visual experience, the authors repeatedly^{viii} claim, is constituted by our implicit knowledge of both kinds of sensorimotor contingency.

Noë (2004) offers, as a more general implication of the sensorimotor approach, the idea that:

"what determines phenomenology is not neural activity set up by stimulation as such, but the way the neural activity is embedded in a sensorimotor dynamic"

Noë (2004) p.227

That this is not merely a claim of contingent causal connection is suggested by the passage that then follows:

"Experience is not caused by and realized in the brain, although it depends causally on the brain. Experience is realized in the active life of the skillful animal. A neuroscience of perceptual consciousness must be an enactive neuroscience—that is, a neuroscience of embodied activity, rather than a neuroscience of brain activity"

Noë (2004) p. 226

If experience is indeed "realized in the active life of the skilful animal", the connection between that skilful life and experience is intimate indeed. For the skilful life of embodied action emerges as quite literally *part of* what makes the experience the way it is.

For my own part (though I shall not try to argue for this here) I am not yet convinced that the case has thus been made for any fully general version of this apparently constitutive claim. The sense in which knowledge of sensorimotor contingencies might here be said to 'constitute' our knowledge of e.g. the visual appearance of straight horizontal lines seems rather to be the sense in which brains like ours might be said to encode information about (visual) straightness and horizontality *by means of* encoding information about the various (both intrinsic and body-relative) signature contingencies. But that sounds like a claim about how certain kinds of knowledge may in fact become available to (the brains of) certain kinds of embodied beings: it thus looks more like a claim about implementation than about the conceptual intertwining of features of embodiment and of mind^{ix}.

There exist, however, a special range of cases, concerning specific kinds of conscious experience, for which a strongly constitutive claim seems quite persuasive. We can introduce them by means of a simple thought experiment.

Try to imagine a creature whose conscious experience presents them with an upside down world, but whose motor routines are so neatly tweaked and tuned that their physical engagements with the world always go off without a hitch. Imagine, moreover, that these beings are so familiar with their own motoric eloquence that they are never surprised that their actions work out. Imagine too that all their episodes of planning and imagination have come to be as well integrated with motoric action as our own, enabling them, for example, to plan and execute complex climbs on mountainsides and indoor training walls, and whatever else you would accept as proof of some proper inter-animation between conscious reason and successful action. Now ask yourself: can you *really* imagine that these beings experience their world as 'upside down'?

Skill-based accounts of conscious perception provide a powerful framework in which to press a negative response. At the heart of such approaches is the simple but compelling idea that *in certain cases* the way we consciously perceive the world is intimately, rather than merely contingently, tied up with routines for (or behavioural dispositions towards) engaging the world by deed and action. One obvious advantage of appealing to behavioural dispositions rather than actual actions is to allow a contingently inactive (perhaps paralyzed or temporarily incapacitated) agent to enjoy the same range of sensory experiences. Such models remain motocentric in a sufficiently strong sense, since they make (certain) conscious contents conceptually inextricable from the potential (which need not be actualized) for embodied action.

Perhaps, then, there is a distinctive range of cases for which something like a motocentric constitutive connection can be made plausible. In particular, such a connection seems especially plausible for the case of spatial content and spatial awareness. Thus Mandik (1999) argues for what he terms the 'behavioral constituency of perceptual space'. This is the idea that our egocentric experience of space is conceptually intertwined with our possession of various bodies of behavioural know-how. For

example, following remarks by Pitcher (1970), Mandik comments that:

"We just cannot imagine hearing a sound coming from a particular direction while simultaneously not knowing how to orient (or point, or walk..) toward the direction from which the sound seems to originate"
Mandik (1999) p. 53

Similar intuitions are pumped in Evans (1985), Taylor (1979) and Grush (1998). Evans, for example, writes that:

"Auditory input, or rather the complex property of auditory input which codes the direction of the sound, acquires a spatial *content* for an organism by being linked with behavioural output"
Evans (1985) p.385

Grush goes on to develop a lovely example concerning the perception of a sound as pulsating. The claim is that:

"part of the normal content of pulsatingness, for us, is that it is something with which we can co-ordinate a number of sensorimotor skills" Grush (1998) para. 21

Thus suppose we hear the sound of a siren as pulsating. Then that perception, Grush argues, poises us to exercise a battery of skills. We might wave a hand, tap a finger, or nod our head in time with the pulses. The total failure of an embodied agent to be able to bring any such skills to bear is, Grush argues, incompatible with the idea that that agent actually perceives the sound as pulsating (though she may know it to be pulsating by some other means). Intrinsic to the perceptual auditory content then, is something that puts that content in touch with dispositions towards various kinds of embodied actions. It is a natural (though admittedly additional) step to suggest that what it is to perceive the sound as pulsating just *is* for such links to be in place.

As an aside, it might be thought that work on the so-called 'dual visual systems hypothesis' (Milner and Goodale (1995)) puts pressure on the idea of a constitutive motocentric link for any form of visual consciousness. This is because such work depicts conscious vision as depending on a set of neural resources that operate semi-independently of those that support fine visuo-motor control. But this, I think (and see Clark (1999) (2001)) is a mistake. For all the motocentric model requires is that the kind of spatial contents available to conscious thinking and reasoning are conceptually bound up with an agent's dispositions to engage a spatially presented world. If the *success* of those engagements also requires the proper operation of multiple other subsystems (e.g. the dorsal stream), this fact is additional to (perhaps even orthogonal to) the constitutive claim itself.

I shall not attempt to further pursue these issues here, since the point was merely to illustrate a specific range of cases, concerning spatial awareness, where a genuinely constitutive connection between motoric dispositions and conscious contents seems suggested by reflection on our ordinary thought and practice. Although this is indeed a limited case, it is worth flagging the possibility that some kind of broadly skill-based account may apply much more generally, perhaps covering all cases of conscious awareness and 'qualia'. Such an extension is explicitly argued for in Noë (2004), where sensorimotor considerations are invoked for experiences of shape, color and perceptual contents of all stripes. Such a general extension gains in plausibility, however, if we drop the requirement that the skills in question be essentially or crucially *sensorimotor* in nature, and allow instead that many will concern passive acts of recognition, of matching to samples, etc. (For a nice treatment of appearances of colour couched in these terms, where the emphasis is on 'sifting, sorting, tracking and judging' rather than on motor engagements per se, see Pettit 2003).

Finally, I should briefly mention a third, and very different way of arguing for CCC. This route appeals to considerations of the role of the body in structuring human concepts. The *locus classicus* here is Lakoff and Johnson's (1980, 1999) work on the role of body-based metaphors in human thought and reason. Many of our basic concepts, they argue, are quite evidently body-based: concepts like front and back, up and down, inside and outside:

“If all beings on the planet were uniform stationary spheres floating in some medium and perceiving equally in all directions, they would have no concepts of *front* and *back*”
Lakoff and Johnson (1999) p.34

But these basic concepts, they go on to argue, end up structuring our understandings (and our inferences) in more rarefied domains. Happiness and sadness, to take the standard example, are conceived in terms of upness and downness. The specifics of human embodiment thus shape the basic concepts that in turn inform (so it is argued) all the rest. Perhaps then:

Organisms that didn't have bodies like our own would develop other metaphors to characterize happiness and sadness. *Happy* and *sad* would be structured in other ways and would thus assume different meanings^x.
Shapiro (2004) p.201

The common upshot of all these arguments, then, is a kind of principled body- or action-centrism, according to which the presence of humanlike minds depends quite directly^{xi} upon either the possession of a humanlike body^{xii} or (more weakly but still significantly) on the disposition to engage in certain distinctive patterns of embodied action.

3. Extended Functionalism

It is revealing, I think, that Shapiro's spirited defense of profound bodily^{xiii} involvement in the mental comes in the larger context of

a series of arguments aimed at a different, logically independent but thematically related, target. That target was the thesis of multiple realizability: a staple of non-reductionist Philosophy of Mind ever since the heady days of early Machine Functionalism. At about that time, the notion that minds like ours might be directly identified with their specific *neural* underpinnings was widely cast as a kind of unacceptable meat or species chauvinism, and replaced by the identification of mind as a functional kind, capable in principle of being realized by many different physical substrates (Putnam (1975)— see also Putnam (1960), (1967)). In this new regime, mindware stood to neural hardware as software stood to the physical device. Just as the same software could turn up on different bedrock machines (as when a PC and a Mac both run Windows), so the same kinds of mind might, it was supposed, turn up in various kinds of material form. What mattered was not the bedrock physical forms so much as the abstract patterns of input-to-internal-state- transitions- to-output that the material structures were able to support. Sameness at this rather abstract level was meant to guarantee sameness at the mental level. Or at any rate, any remaining slack was to be taken up by rather arcane details of history and/or distal environmental embedding. As far as the machinery of mind itself was concerned, functional identity fully fixed any contribution to mentality.

Shapiro's appeal to work in embodied, embedded cognitive science depicts it as in spirit rather inimical^{xiv} to the platform-neutral machine functionalist model of mind. But there is a way of understanding the embodied, embedded approach that sees it as extending, rather than undermining, a broadly functionalist story. To bring this possibility into focus, we must next sample the kinds of argument that favour what we earlier dubbed EF— for Extended Functionalism.

Arguments in favour of EF appeal mainly, if not exclusively, to the *computational role* played by certain kinds of non-neural events and processes in online problem-solving. Consider, to pursue an example cited earlier in the text, Ballard et al's (1997) account of the use of 'deictic pointers' in a block-copying task. In this task, a

subject is given a target pattern of colored blocks that they must copy by moving similar blocks from a reserve area to a new workspace. Using the spare blocks in the reserve area, the subject must recreate the pattern by moving one block at a time from the reserve to the new version they are creating. The task is performed using mouse clicks and a monitor, and the subject's eye motions are constantly tracked. What Ballard et al found was that repeated rapid saccades to the model were used in the performance of the task: many more than you might expect. For example, the model is consulted both before and after picking up a block, suggesting that when glancing at the model, the subject stores only one piece of information: either the color or the position of the next block to be copied.

To test this hypothesis, Ballard et al used a computer program to alter the color of a block while the subject was looking elsewhere. For most of these interventions, subjects did not notice the changes even for blocks and locations that had been visited many times before, or that were the focus of the current action. The explanation was that when glancing at the model, the subject stores only one piece of information: either the color or the position of the next block to be copied (not both). In other words, even when repeated saccades are made to the same site, very minimal information is retained. Instead, repeated fixations provide specific items of information 'just in time' for use. The conclusion from this is that:

"In the block-copying paradigm... fixation appears to be tightly linked to the underlying processes by marking the location at which information (e.g., color, relative location) is to be acquired, or the location that specifies the target of the hand movement (picking up, putting down). Thus fixation can be seen as binding the value of the variable currently relevant for the task"

Ballard et al (1997) p 734

Two morals matter for the story at hand. The first is that visual fixation is here playing an identifiable computational role. As the

authors (op cit p.725) comment “changing gaze is analogous to changing the memory reference in a silicon computer”. The second is that repeated saccades to the physical model thus allow the subject to deploy what Ballard et al dub ‘minimal memory strategies’ to solve the problem. The idea is that the brain creates its programs so as to minimize the amount of working memory that is required, and that eye motions are recruited to place a new piece of information into memory. Indeed, by altering the task demands, Ballard et al were also able to systematically alter the particular mixes of biological memory and active, embodied retrieval recruited to solve different versions of the problem. They conclude that, in this kind of task at least, “eye movements, head movements, and memory load trade off against each other in a flexible way” (op cit p.732)

The Ballard et al model is a clean example of an extended functional approach. It analyses a cognitive task as a sequence of less intelligent sub-tasks, using recognizable computational and information-processing concepts, but applies those concepts within a larger organizational whole. It recognizes the profound contributions that embodiment and environmental embedding make to the solution of the problem, and displays those contributions rather clearly, by identifying the abstract role of specific (both gross-bodily and neural) operations in real-time performance of the task. The authors are fully aware of this, commenting that their model “strongly suggests a functional view of visual computation where different operations are applied at different stages during a complex task” (op cit p. 735). As a result, a Ballard-style approach is able:

“To combine the concept that looking is a form of doing with the claim that vision is computation [integrating the two points by] introducing the idea that eye movements constitute a form of deictic coding...that allow perceivers to exploit the world as a kind of external storage device” Wilson (2004) p.176–177

Bodily actions are thus part of the means by which certain computational and representational operations are implemented. But what makes the cognitive process the one that it is is simply its functional profile (the set of state transitions mediating input and output). The difference is just that this functional profile belongs not to the neural system and its inputs and outputs alone, but to the whole embodied system located in the world. We thus have a version of what Block (1990, p.70) calls 'long-arm functionalism' viz functional roles that "reach out into the world of things".

A second example of extended functionalism is Clark and Chalmers (1998) account of the 'extended mind'. The point of that argument (which I shall not rehearse again today) was to show that external memory traces may sometimes be poised for the control of action in very much the same kind of way as internal memory traces, yielding an extended supervenience base for dispositional beliefs. The claim here was not, implausibly, that an external, passive, encoding might somehow behave exactly like the fluid, automatically responsive resources of internal biological memory. Rather, it was that external encodings were, under certain circumstances, capable of becoming so deeply integrated into online strategies of reasoning and recall as to be only artificially distinguished from proper parts of the cognitive engine itself. The argument thus echoes Ballard et al's in depicting a larger integrated system as the extended machinery whose computationally salient states, properties and transformations are supposed to explain specific problem-solving performances.

Or consider an accountant who is extremely good at dealing with long tables of figures. Over the years, she has learnt how to solve specific classes of accounting problem by rapidly scanning the columns, copying some numbers onto a paper scratchpad, then looking to and from those numbers (carefully arrayed on the page) back to the columns of figures. The accountant (let's call her Ada) does this at lightning speed, and by deploying a variety of minimal memory strategies. Instead of committing complex dependencies to biological memory, Ada follows trails through the numbers,

creating external traces every time an intermediate result is obtained. These traces are in turn visited and re-visited on a ‘just in time, need to know’ basis, moving specific items of information into and out of short term bio-memory in much the same way as a serial computer shifts information to and from the central registers in the course of carrying out some computation. This process may again be best analysed in extended functional terms, as a set of problem-solving state-transitions whose implementation happens to involve a distributed combination of biological memory, motor actions, external symbolic storage, and just-in-time perceptual access.

R.Wilson’s (2004) notions of ‘exploitative representation’ and ‘wide computation’ capture the key features nicely. Exploitative representation occurs when a sub-system gets by without encoding some piece of information, in virtue of its ability to track that information in some other way. Wilson gives the example of an odometer that keeps track of how many miles a car has traveled not by first counting wheel rotations then multiplying according to the assumption that each rotation = x meters, but by being built so as to record x meters every time a rotation occurs:

“ In the first case it encodes a representational assumption and uses this to compute its output. In the second it contains no such encoding but instead uses an existing relationship between its structure and the structure of the world” Wilson (2004) p.163

Wilson’s descriptions and central examples can make it seem as if exploitative representation is all about achieving success without representations at all, at least in any robust sense of representation. But this need not be so. Another, very pertinent, range of cases would be those in which a sub-system does not contain within itself a persisting encoding of certain things, but instead leaves those encodings in the world (or in some other sub-system to which it has access). In this kind of case, the larger system may well contain genuine symbolic encodings, consistent with the sub-system merely exploiting (rather than

recapitulating) these during problem-solving activity. The case of Ada, described above, would be a case in point. Ada's biological brain does not create and maintain persistent encodings of every figure she generates and offloads onto the page, though it may very well create and maintain persistent encodings of several other key features (for example, some kind of running approximation that acts to check for gross errors). In much the same way as Ballard's block-puzzlers, Ada's biological brain may thus, via the crucial bridging capacities of available embodied action, key its own internal representational and internal computational strategies to the reliable presence of the external pen-and-paper buffer. Even robustly representational inner goings-on may thus count as exploitative insofar as they merely form one part of a larger, well-balanced process whose cumulative set of state-transitions solves the problem. In this way:

“explicit symbolic structures in a cognizers environment...together with explicit symbolic structures in its head [may] constitute the cognitive system relevant for performing some given task” Wilson (2004) p.184

The sense of constitution here in play is not, however, the strong sense of "being conceptually intertwined with" but the weaker (though still interesting and important) sense of "as it happens being part of the apparatus that implements".

The use of various forms of exploitative representation immediately yields a vision of what Wilson dubs 'wide computationalism', according to which "at least some of the computational systems that drive cognition reach beyond the limits of the organismic boundary" (op cit p.165). Wide computationalism, stressing at it does the many interactive processes that span brain, body and world, is also intrinsically dynamics-friendly. Many of the internal representational states invoked will be fleeting, generated on-the-spot, delicately keyed to making the most of other closely coupled internal and external resources. Extended functionalism is thus not in any way committed to their being static symbols in the head. On the

contrary, it invites us to locate static, classical-looking symbol structures where they belong: out in the world, but making a deep and abiding contribution to online thought and reason nonetheless.

Once again, the point is not to defend this specific suggestion, so much as to locate it in the wider space of options. The key features, for these purposes, are that extended functional systems may include motor behaviours as processing devices, and environmental structures as storage and encoding devices^{xv}. Such bodily and worldly elements emerge as genuine parts of extended computational regimes, and are apt for formal description in both dynamical and information-processing terms^{xvi}. The larger systems thus constituted are, as Wilson insists, unified computational wholes such that “the resulting mind-world computational system itself, and not just the part of it inside the head, is genuinely cognitive” (op cit p.167). Insofar as such a thesis is correct, the cognitive scientist or philosopher of mind who chooses to treat the brain and central nervous system alone as the mechanistic supervenience base for mental states is rather like a neuroscientist who insists that neuroscience proper should not be concerned with the hippocampus or the cerebellum, because (they think) *all the real cognizing goes on in the cortex*, even if (they concede) those other structures sometimes play a role in the transmission and routing of information and control.

Extended functionalists^{xvii} thus reject the image of mind as a kind of input-output sandwich with cognition as the filling (for this picture, and more arguments for its rejection, see Hurley (1998)). Instead, we confront an image of the local mechanisms of human cognition quite literally extending out into the world. But just as the traditional machine functionalist was interested in neural goings-on as the contingent means by which human beings manage to implement a particular functional organization (since it is the functional organization itself that then bears some constitutive relation to our mental states) so the extended functionalist sees the body as the contingent bridging mechanism that allows us to implement some larger functional organization,

that in turn may bear some conceptual or constitutive relation to our mental states.

We are now in a position to bring out the potential tension between EF and CCC: a tension somewhat obscured by the common use of the notion of embodiment. For while embodiment and embedding play a crucial role in each kind of story, the nature of that role is often rather different. Thus, the accounts mentioned in section 2 seem to depict bodily form and/or knowledge of sensorimotor contingencies as elements that make a direct and constitutive contribution to human thought and reason. While the accounts mentioned in section 3 depict bodily action as providing a contingent bridge for the emergence and maintenance of larger functional wholes spanning a wider computational, dynamical, and representational nexus^{xviii}. What counts, for mind, cognition, and mental states, according to these latter models, is that overall processing economy itself.

To dramatize the difference, recall Shapiro's opposition to the idea that "the same kind of mind can exist in bodies with very distinct properties" (Shapiro (2004) p.175). On the basis of the kinds of evidence described in section 2 above, Shapiro rejects the idea that "snakelike organisms and creatures of science fiction" (op cit 174) might share our kind of mind. If the theorists of embodied cognition are correct, Shapiro suggests, Body Neutrality (the idea that "characteristics of bodies make no difference to the kind of mind one possesses" (op cit p.175)) is false.

It should be clear, however, that something has here gone by a little too swiftly. For imagine now a case in which we have two intelligent beings. One of them is a snake-like creature lying on top of an advanced touch-screen like environment. In this flat-screen setting every little wriggle of the snake can cause specific external symbolic tokens to appear elsewhere on the screen: tokens that are themselves apt for perceptual uptake (perhaps via a kind of Braille). The snake-being (call it Adder) uses this set-up, let us suppose, to implement the same extended accounting

algorithm as Ada in our earlier example. According to EF there is no reason to suppose (at least from anything we have said so far) that the accounting–relevant cognitive states of Ada and Adder need differ in any respect. Each implements the same extended computational process, and even (we may suppose) divides the biological and non–biological contributions in the same way.

More radically, but still consistent with EF, we may even imagine that there are differences at the level of what gets done where. Enter Odder. Odder performs certain computations internally that Ada and Adder both perform using action and perception routines in the non–biological arena. Here too, the EF theorist is at liberty to believe that the very same cognitive and mental states might be being implemented, with nothing distinguishing the cases apart from some non–essential matters of location. Just as, for the standard approach, we need not care (within sensible limits) exactly where *within* the brain a given operation is performed, so too (it might be urged) we should not care whether, in some extended computational process, a certain operation occurs inside or outside some particular membrane or metabolic boundary.

Odder is by no means a duplicate in the sensory arena, and the distribution of operations varies widely between Odder, Ada, and Adder. Moreover, the distribution varies in ways that criss–cross the perceptual/non–perceptual divide. Yet there is a strong (full-bloodedly) functionalist intuition that *considered as a form of 'accounting intelligence', Odder, Adder and Ada form a computational equivalence class.*

Returning to Shapiro's argument we can now see that, looked at from the perspective of EF, it mistakenly moves from:

- (1). Bodily structures and action routines can be key participants in extended information–processing routines.

to

(2). Bodily structures and action routines play a special role such that *sameness of mental state requires sameness of bodily structure and action routine*.

The spirit of EF, it seems to me, is precisely to reject (2) by stressing overall systemic integrity at the expense of any special role for the contributions of body, brain, or world. Fans of CCC, by contrast, may think that EF fails to do justice to what matters most about the body and action, and that they are simply repeating an old functionalist mistake in a new (extended, distributed) setting.

4. Remedies for the Flesh?

There is, I have argued, a potential tension at the very heart of the program so easily (so unitarily) glossed as the study of “embodied, embedded cognition”. It is the tension between an extended, dynamics-friendly, version of good old fashioned functionalism, and something more fundamentally fleshy: the idea that features of the body (including characteristic patterns of embodied action) make a direct and in some sense profound or constitutive contribution to mind and mentality. My goal in the present treatment has been simply to expose this tension in as clear a way as I can. I’d like to end, however, with a brief sketch of some possible steps towards a resolution.

One quick way to relieve the tension would be to argue that the body makes a direct and constitutive contribution to conscious (or perhaps just to conscious perceptual) experience but not to cognition considered more generally. Thus Ada, Adder and Odder may indeed all implement the very same accounting algorithm, but the differences in embodiment will (on this account) necessarily make a difference at the level of their conscious experience. The wild card in this whole debate is thus our old friend phenomenal experience itself. Thus consider Shapiro’s observation that:

“The instructions by which the human brain computes relative depth do not work in creatures with eye configurations other than those in a human being. This is the sense in which depth perception is embodied. The procedures by which human beings perceive depth— a fact about human psychology— are contingent on a fact about human bodies” Shapiro (2004) p.188

Recall that from facts such as these, Shapiro concludes that “human vision needs a human body” (op cit p.189). Such a claim is, however, importantly ambiguous. It might mean only (1) that the brain’s own algorithms factor in the opportunities provided by bodily structures and action. It might mean (2) that being able to make the kinds of gross visual discrimination that we can make *requires* having exactly the same kind of body (in respect of eye configuration at least) as we do. But this claim is surely false, since an alternative distribution of the very same information processing steps, in some differently-brained and differently bodied being, must be capable of implementing that very same algorithm (as EF is frequently at pains to point out)^{xix}. Or it might mean (3) that any such alternative implementation need not preserve the qualitative feel of human depth perception: a qualitative feel that is somehow tied not to the running of a certain abstract algorithm but to the use of two eyes located a certain distance apart.

The quickest reconciliation, then, is to treat the body or sensory apparatus as itself making some special kind of contribution, one that cannot help but impact (in non-trivial ways) certain qualitative aspects of our mental life. This is probably the best way to understand Noë 's assertion that "the character of our experience depends on...idiosyncratic aspects of our sensory implementation" (op cit p.26). If you think that the sensory implementation plays a unique role that contributes directly to experiential content, you may very well think that every difference in implementation makes a real (though perhaps vanishingly small) difference to the felt nature of the experience itself. In this way, you get both to be a kind of functionalist

(arguing that it is the role of the physical structures in mediating patterns of sensorimotor contingency that matters) *and* to assert that (for perceptual experience at least) every difference in certain aspects of physical implementation makes a difference.

It is by no means obvious, however, that we can stably reconcile any recognizable form of functionalism with such *full-and-principled* sensitivity to all the details of a being's embodiment and/or sensory apparatus. For a broadly functional (or even just a broadly computational and representational) view of the underpinnings of perceptual experience demands, it seems to me, that it be in principle possible that two beings could be different in respect of gross sensory apparatus and embodiment and yet, courtesy perhaps of compensatory differences in key aspects of downstream processing, end up realizing the same set of experience-determining functionally specified state-transitions. Noë (2004) (and also O'Regan and Noë (2001)) seem to leave no room for this even as a bare possibility. Noë is explicit that "to see *as we do*, you must...have a sensory organ and a body like ours" (Noë (2004) p.112, italics in original)^{xx}.

Perhaps this is right, and experience is non-trivially permeated by the full details of biological embodiment. My own view (see e.g. Clark (1999) (2001): see also Jacob and Jeannerod (2003)) is that this is unlikely to be true. Experience, in presenting us with a world fit to engage by action and by reason, need be sensitive only to certain aspects or features of the sensorimotor contingencies our embodiment dictates. Broader skill-based approaches have more room for maneuver here, as we saw briefly in section 2 above. By simply identifying experiences with implicit knowledge of the full suite of contingencies defined at the sensorimotor surfaces, the strong sensorimotor account leaves no room for compensatory downstream adjustments to yield identical experiences despite surface dissimilarities^{xxi}. Nor does it leave room for small differences at the sensorimotor surfaces to be such as to make no experiential difference, courtesy of failing to deliver any *salient* differences in signals to downstream processors. Perhaps, that is to say, downstream processing provides a kind of

grid relative to which certain differences at the level of the sensory inputs (and associated contingencies) simply fail to *make* a difference. But whatever the cognitive scientific niceties, the point for present purposes is just that no global reconciliation between EF and the strong sensorimotor model looks likely, since the latter depicts the sensorimotor dynamics alone as fixing, with extreme sensitivity, the nature of our perceptual experience. This forecloses the possibility of the same experience being brought about in some other sensorimotor context in virtue of the details of a larger functional organization (or, contrariwise, of a different perceptual experience being brought about, despite sameness of sensorimotor contingencies, in virtue of some larger functional organization).

A related worry threatens at least the strongest versions of Lakoff and Johnson's claims concerning the tight links between forms of embodiment and basic conceptual repertoires. For what embodied experience actually delivers as the baseline for learning and metaphorical thought surely depends on some complex mixture of bodily form, environmental structure and (possibly innate) downstream internal processing. Here too, compensatory adjustments in either of the two non-bodily arenas look likely to make available forms of thought and reason that are not tethered in any simple way to the gross bodily bedrock.

For these reasons, I think it would be unwise to opt for a global reconciliation according to which bodily form and sensorimotor patternings play a profound or constitutive role in the determination of all forms of conscious perceptual experience. Instead, we should try, on a case-by-case basis, to understand the body as playing an enabling computational role: one that selectively impacts both conscious and non-conscious cognitive strategies. In this way we may hope to explain, for a variety of specific cases, just why and how the body matters, without making the body matter mysteriously or simply in and of itself.

One obvious place to start is by observing that for many information processing problems, there will be an elegant,

representationally low-cost, solution that is in several important senses relative to the gross physical properties of the implementing apparatus. To borrow an example that I have used before, an industrial solution to the problem of fitting together small highly engineered parts may be built around the provision of flexible rubber mountings for the assembler arms. The 'bodily' structure here reshapes the information processing required. Subtract the rubber mountings and the best solution involves the repeated computation of multiple visual feedback loops to guide repeated attempts. But with the rubber mountings in place, the system gives along two spatial axes, and even roughly visually guided parts fall into place "just as if millions of tiny feedback adjustments to a rigid system were being continuously computed" (Michie and Johnston 1984 p.95).

Similar effects sometimes flow from the concrete details of sensor placement. A system with a certain spatial distribution of sensors for heat or light will not need to deploy multiple steps of inference to determine whether certain signature patterns are present or absent. Moreover, the fixed relations between bodily-mounted sensors obviate the need to constantly determine how input at point X relates to input at point Y. Such relations may be either constant (as between two fixed eyes) or else vary systematically (where X and Y are independently controllable or moveable, as in the case of the left and right index fingers). In either case, the properties of the body keep the sensory inputs in a certain kind of alignment, and this invariant can be simply assumed (rather than explicitly represented) by the algorithms that use the sensory inputs as sources of problem-solving information.

The body is also the point at which willed action, if successful, first impacts the wider world. This sounds trivial, but is actually profoundly important. When conjoined to the observation that, in the typical human case, these points of willed action include all our voluntary sensor movements, it yields the intuitive understanding of the body as the common and persisting locus of sensing and action. Extensive work on the technologies of telepresence^{xxii} suggests that the human sense of presence, of

being at a certain place in space, is fully determined by our ability to enter into closed loop interactions, in which willed sensor motions yield new sensory inputs, and by our ability to act upon at least some of the items thus falling within sensory range. The phenomenological sense in which we typically 'inhabit' (rather than painstakingly control) our own bodies is determined by the fluency with which we control the sensors and actuators in ways that enable successful interaction to occur.

Insofar as our general sense of presence in the world is so determined, it should come as no surprise that our sense of being oriented in the world (e.g. of seeing the world right way up or upside down) is also tightly synched to our capacities for (or better, our dispositions towards) successful fluid sensorimotor engagements. The most convincing cases for constitutive *motor* involvement in perceptual experience were all seen to fall within this limited class: the class of phenomenal contents relating to orientation, location and presence.

Finally, the body, by being the immediate locus of willed action, is also the gateway to intelligent offloading. The body is the primary tool for the intelligent use of environmental structure (see Kirsh (1995)) and acts as the mobile bridge that allows us to use the external world in ways that simplify and transform internal problem-solving. The body is thus the go-between linking these two different (internal and external) sets of key information-processing resources. The body's role in such cases is that of an instrument enabling the emergence of a new kind of information-processing organization. This role may, without too much exaggeration, be likened to that of the corpus callosum. Both are key physical structures whose cognitive role is in part to allow distinct sets of resources to engage in highly integrated forms of problem-solving activity.

In these closing comments I have been speaking as if the body is, just as it happens, the locus of willed action, the point of sensorimotor confluence, the gateway to intelligent offloading, and the stable (though not permanently fixed) platform whose features

and relations can be relied upon in the computation of certain information-processing solutions. But in fact, I am inclined to go further and to assert not just that this is what the body does, but that this (or some list quite like it) is what, at least for all cognitive scientific purposes, the body *is*. I am inclined, that is, to identify the body with whatever plays these roles in a unified information-processing economy.

The immediate upshot of this is that the body, insofar as it is cognitively significant, turns out to be itself defined by a certain complex functional role. Notice also that nothing here requires a single persisting body in ordinary 3-space. Instead, there could be genuine but scattered forms of embodiment, embodiment in virtual or mixed realities, and multiple embodiments for a single intelligence (for more on these topics, see Clark (2003) (In Press) (Jenann (ms) (forthcoming))). It is merely a contingent (and increasingly negotiable) fact about human embodiment that the body is both the metabolic centre *and* the bridge to sensory presence and intelligent action. In exotic cases^{xxiii}, the metabolic center is detached from the more cognitively important loci for sensing, acting and intelligent offloading: loci that collectively determine our sense of presence in the world.

A deeper resolution of the apparent tension between EF and even the more sweeping forms of CCC is now in sight. For the cognitive importance of the body, if this is on the right track, is fully exhausted by its ability to play a certain functional role in an intelligent organization. The distinctiveness and importance of this role is what explains the intuition that the body makes a special and quite pervasive cognitive contribution. But because it *is* nothing but a complex functional role, there is nothing cognitively significant about the bodily contribution that is not fully captured by reflection upon its several (and potentially separable) computational and information-processing contributions. The body *is* special. But we should understand its specialness through the familiar lens of our best information-processing models of mind and cognition.

6. Conclusions: The Body's Balancing Act

There is a certain tension, or so I have argued, between two strands of thought prominent in the recent literature on the embodied, embedded mind. One of those strands depicts the body as intrinsically special, and the details of a creature's embodiment as a constitutive, or at any rate profound and abiding, constraint on the nature of its mind. The other depicts the body as just 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. This is a kind of extended functionalism (now with an even broader canvas for multiple realizability than ever before).

One possible resolution of this tension is to display the body as (for all cognitive purposes) nothing but the item, or items, that play a certain complex functional role in an information-processing economy. Within such an economy, mental sameness is indeed determined by the overall balance achieved using neural, bodily and environmental resources. The body plays a special (but plainly instrumental) role in determining and stabilizing this balance and as such it is a key player on the cognitive stage.

There is, however, an important sub-class of cases in which our dispositions for embodied action may plausibly be depicted as playing a strongly constitutive (not merely instrumental) role. These are the cases concerning space, orientation and presence. This class may yet be extended by new skill-based models that look beyond the sensorimotor surfaces, discerning a constitutive involvement for various acts of recognizing, grouping and tracking. We have found no reason, however, to suppose that in any of these cases mental sameness requires gross bodily sameness.

* Thanks to Michael Wheeler, Jenann Ismael, David Chalmers, Mark Rowlands, and Rob Wilson for useful comments during the early stages of this project.

References

Adams, F and Aizawa, K (2001) "The Bounds of Cognition" *Philosophical Psychology* 14:1 p.43-64

Ballard, D.H., M.M. Hayhoe, P.K. Pook, and R.P.N. Rao (1997) "Deictic codes for the embodiment of cognition," *Behavioral and Brain Sciences* 20, 4, 723–767.

Block, N (1990) "Inverted Earth" in J. Tomberlin (ed) *Philosophical Perspectives 4: Action Theory and Philosophy of Mind* (Ridgeview Press, CA)

Churchland, P.M. (2005) "Functionalism at Forty: A Critical Retrospective" *Journal of Philosophy*. Vol CII: 1: 33-50

Clark, A (1997) *Being There: Putting Brain, Body and World Together Again* (MIT Press, Camb. MA).

Clark, A (1999) "Visual Awareness and Visuomotor Action" *Journal Of Consciousness Studies* 6:11-12:1-18

Clark, A (2001) "Visual Experience and Motor Action: Are The Bonds Too Tight?" *Philosophical Review* 110:4: 495-519

Clark, A (2003) *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence* (Oxford University Press, NY).

Clark, A (In Press) "Reinventing Ourselves: The Plasticity of Embodiment, Sensing, and Mind". *Journal of Philosophy and Medicine*,

Clark, A. and Chalmers, D. (1998). "The Extended Mind." *Analysis* 58: 10-23 (Reprinted in (P. Grim, ed) *The Philosopher's Annual*, vol XXI, 1998).

Clark, A and Toribio, J (2001) "Sensorimotor Chauvinism?: Commentary on O'Regan, J.K., and Noë, A. "A sensorimotor approach to vision and visual consciousness" *Behavioral and Brain Sciences* 24: 5: 979-980

- Damaio, A (1994) *Descartes' Error* (Grosset/Putnam, NY)
- Damasio, A (1999) *The Feeling of What Happens* (Harcourt, Brace and Co, NY)
- Dennett, D (1981) "Where am I?" in D. Dennett, *Brainstorms* (Harvester Press, Sussex)
- Dennett, D (1996) *Kinds of Minds* (Basic Books, NY)
- Donald, M (1991) *Origins of the Modern Mind* (Harvard University Press, Cambridge, MA)
- Evans, G (1985) "Molyneux's Question" in G. Evans, *The Collected Papers of Gareth Evans* (Oxford University Press, London)
- Gallagher, S (2005) *How the Body Shapes the Mind* (Oxford University Press, Oxford)
- Grush, R (1998) "Skill and Spatial Content" *Electronic Journal of Analytic Philosophy (EJAP)* Volume 6 (<http://ejap.louisiana.edu/EJAP/1998>).
- Grush, R (2003) "In Defense of some 'Cartesian' Assumptions about the Brain and Its Operation" *Biology and Philosophy* 18:53-93
- Haugeland, J (1998) "Mind Embodied and Embedded" in J. Haugeland *Having Thought: Essays in the Metaphysics of Mind* (Harvard University Press, Cambridge, MA) 207-240, (this originally appeared in *Acta Philosophica Fennica* 58 (1995) p.233-267 (a special issue on Mind and Cognition edited by L. Haaparanta and S. Heinamaa)).
- Houng, Y and Ho, J (eds) (1995) *Mind and Cognition* (Academia Sinica, Taipei)
- Hurley, S (1998) *Consciousness in Action* (Harvard University Press, Cambridge, MA)
- Hutchins, E (1995) *Cognition in the Wild* (MIT Press, Camb. MA)
- Ismael, J (ms) "The self, the body, and space"
- Ismael, J (forthcoming) *The Situated Self* (Oxford University Press, Oxford)

- Jacob, P and Jeannerod, M (2003) *Ways of Seeing: The Scope and Limits of Visual Cognition* (Oxford University Press, Oxford)
- Kirsh, D (1995) The Intelligent Use of Space. *Artificial Intelligence* 73(1-2): 31-68
- Lakoff, G, and Johnson, M (1980) *Metaphors We Live By* (University of Chicago Press, Chicago)
- Lakoff, G and Johnson, M (1999) *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought* (Basic Books, NY)
- Mandik, P (1999) "Qualia, Space, and Control" *Philosophical Psychology* 12:1: 47-60
- Michie, D and Johnson, R (1984) *The Creative Computer* (Penguin, NY)
- Milner, A and Goodale, M (1995) *The Visual Brain in Action* (Oxford University Press, Oxford, UK)
- Noë, A (2004) *Action in Perception* (MIT Press, Camb. MA)
- O'Regan, J.K., and Noë, A. (2001) "A sensorimotor approach to vision and visual consciousness" *Behavioral and Brain Sciences* 24: 5: 939-973
- Parfit, D (1984) *Reasons and Persons* (Oxford University Press, Oxford),
- Pettit, P (2003) "Looks as Powers" *Philosophical Issues* 13:221-252
- Pitcher, G (1970) *A Theory of Perception* (Princeton Univ Press, Princeton, N.J.)
- Putnam, H (1960) "Minds and Machines" in S. Hook (ed) *Dimensions of Mind* (New York University Press, NY)
- Putnam, H (1967) "Psychological Predicates" in W. Capitan and D. Merrill (eds) *Art, Mind and Religion* (Pittsburgh University Press, Pittsburgh, PA) 37-48
- Putnam, H (1975) "Philosophy and our Mental Life" in H. Putnam (ed) *Mind, Language, and Reality* (Cambridge University Press, Cambridge, UK) 291–303.
- Rowlands, M (1999) *The Body in Mind: Understanding Cognitive Processes* (Cambridge University Press, Cambridge, UK)

Shapiro, L (2004) *The Mind Incarnate* (MIT Press, Camb. MA)

Spencer, J and Schöner, G (In Press) "Bridging the representational gap in the dynamic systems approach to development" *Developmental Science*.

Taylor, G (1979) "The Validity of Transcendental Arguments" *Proceedings of the Aristotelian Society* 79: 151-165

Varela, F., Thompson, E., and Rosch, E (1991) *The Embodied Mind* (MIT Press, Camb. MA)

Wilson, M (2002) "Six views of embodied cognition" *Psychonomic Bulletin and Review* 9:4:625-636

Wilson, R. A. (2004) *Boundaries of the Mind: The Individual in the Fragile Sciences--Cognition* (Cambridge University Press, Cambridge, UK)

ⁱ A very modest sampling might include Thompson, Varela and Rosch (1991), Clark (1997), Haugeland (1998), Hurley (1998), Clark and Chalmers (1998), Rowlands (1999), Lakoff and Johnson (1999), O'Regan and Noë (2001), Shapiro (2004), Wilson (2004), Noë (2004), Gallagher (2005). For some critiques, see Adams and Aizawa (2001), Grush (2003). It should go without saying that many if not all of the key ideas have longer and much more varied pedigrees. For a brief survey of these historical roots, see Clark (1997) chapter 8, section 8.

ⁱⁱ I first heard the paper at an Academia Sinica meeting in Taipei in 1993. It appeared in a volume following that meeting (Yu–Houng Houng and Jih–Ching Ho (eds) (1995). The version quoted in the present text is the one found in Haugeland (1998).

ⁱⁱⁱ Consider, by way of analogy, the familiar claim that memory is (wholly or partially) constitutive of personhood. We can distinguish this interesting, venerable (though probably false) idea from the less interesting observation that memory is often causally implicated in the development of persons. The latter is surely true, while the former is thrown into doubt by the many thought experiments (and indeed real-world cases) in which memory and personhood seem to come apart. For some classic discussion, see Parfit (1984).

^{iv} An exception to this rule is R.A. Wilson (2004)

^v For this reason, Clark and Chalmers (1998) dub their position ‘active externalism’.

^{vi} Shapiro (2004) distinguishes claims concerning body–specificity, claims concerning the role of language and metaphor, and claims concerning tool–based cognitive extension.

^{vii} Shapiro’s larger project is to attempt to undermine the Multiple Realizability Thesis (MRT) according to which the very same mind could exist in many different physical incarnations. The argument proceeds via two steps. First, he argues that MRT implies that we should not in general be able to predict properties of the brain from properties of the mind. Second, he tries to show that such predictions are in fact possible (for example, there are reasons to think that many details of wiring and placement of neural structures are highly constrained by what those structures are required to do), and thus concludes that MRT is false. The details of these arguments will not concern us today. It is perhaps ironic, though, that at least some of the kinds of story that Shapiro appeals to later in the book, as part of his subsequent argument for the importance of specific bodily forms to specific kinds of mind turn out to exemplify a kind of extended functionalism, and thus threaten (or so I shall later suggest) to increase, rather than reduce, the opportunities for multiple realizability of mental states.

^{viii} For example, sensory modalities are said to be "constituted by distinct patterns of sensorimotor contingency" (op cit p.943), as is "the qualitative character of experience" (p.960), "looking round" (p. 968) "perceived color" (footnote 20 p.972) and ultimately perception itself (p.1020). In their reply to Block (p.1015), the authors note that, at least insofar as qualitative experience is concerned, it is not the sensorimotor contingencies themselves that are meant to play the constitutive role but "the perceiver's exercise of mastery of laws of sensorimotor contingency" (p.1015). This is important for distinguishing their view from the hardline behaviorist reading offered by Block, and is the reading assumed throughout the present treatment.

^{ix} It is tempting to think that any conceptual intertwining must actually be of some much more general kind, for example that *some or other* knowledge of sensorimotor contingencies is inevitably implicated in all forms of perceptual experience, but that the very same experience may, in differently embodied beings, be implemented by different bodies of knowledge of sensorimotor contingencies. Nothing in the text, however, seems to suggest such a weakened reading.

^x That there is something potentially problematic about this argument is evident in the tension between the easy use of a common notion of happiness and sadness in the first quoted sentence and the subsequent assertion that happy and sad would then ‘assume different meanings’. But the point, in any case, is simply that arguments stressing the pervasive influence of embodiment on conceptualization look to be arguments for UC, since they assert the ineliminable involvement of bodily details in an account of mental states.

^{xi} Recent work on the complex relations between affect and embodiment (Damasio (1994) (1999)) might seem to establish just such a direct link. But here too we need to proceed with some caution. For the question is not whether or not gross bodily states and processes play some important role in the determination of mental states, but whether *specific details* of embodiment play a *profound or constitutive* role. One test for this is to ask whether a creature lacking a certain kind of body could nonetheless enjoy those very same mental states. This is the question to which Shapiro, Noë and O’Regan, and Lakoff and Johnson (all for very different reasons) all seem to answer ‘no’. Sameness of (at least some) mental states, they believe, requires sameness of key features of embodiment. Damasio’s answer to this question is, however, unclear. The body matters, for Damasio, insofar as low-level bodily responses provide a set of ‘somatic markers’ that afford a kind of compact summary of previous experiences able to impact choice and reasoning in important and often unsuspected ways. Somatic markers thus require there to be some body or other in the loop. But they do not look to require the presence of a body of any specific type.

^{xii} Thus: "For two systems to have the same knowledge of sensorimotor contingencies *all the way down* they will have to have bodies that are identical *all the way down* (at least in relevant respects). For only bodies that are alike in low-level detail can be functionally alike in the relevant ways" (O’Regan and Noë (2001), p.1015)

^{xiii} I use ‘bodily’ here to refer to the gross physical body, excluding the brain and central nervous system.

^{xiv} Inimical to, but not inconsistent with. ST is said to be logically independent of MRT (Multiple Realizability Thesis) since “it is logically possible that a mind could be realized in a number of different kinds of structure, but that all of these structures are contained in similar sorts of bodies (and) it is logically possible that there is only one or a few ways of realizing a humanlike mind but that these few types of realizations can exist in many different sorts of bodies” (Shapiro (2004) p.167). Such concessions make the intended

force of the earlier arguments depicting physical structures as proper parts of psychological processes unclear, though Shapiro does add that he is willing to bet that “if there are but a few ways to realize a humanlike mind, probably there are but a few kinds of bodies that could contain such a mind’ (op cit p.167).

^{xv} Concerning memory, for example, Wilson writes:

“Memory...does not simply stop at the skin but involves engaging with the world through cognitively significant, embodied action...Remembering, on this view, involves exploiting internal, bodily and environmental resources in order to produce some sort of action...”
Wilson (2004) p.191

^{xvi} The idea that dynamical approaches are incompatible with computational and representational ones is increasingly recognized as a mistake, even by those working at the very heart of the dynamical systems movement: see e.g. Spencer and Schöner (In Press).

^{xvii} For example Wilson (2004), Clark (1997), Clark and Chalmers (1998), Dennett (1996), Donald (1991), and Hutchins (1995) all count as extended functionalist treatments according to the present usage.

^{xviii} P.M. Churchland (2005) argues at length that functionalism is a theory that science and philosophy have now outgrown. I am not convinced, however, that the critique is sufficiently sensitive to the very possibility laid out in the present text viz that the most persuasive form of functionalism depicts many of the relevant functional organizations (especially those pertaining to higher-level human cognition) as implemented not by the biological brain alone but by hybrid systems encompassing brain, body, world and action.

^{xix} Thus consider FLICKER. Flicker is a creature with just one eye that moves very rapidly from side to side of its face. Perhaps with some canny tweaks of the neural control and downstream sensory post-processing circuitry, such a being could nonetheless implement precisely the same basic depth perception algorithm as ourselves. The situation would be not unlike the use of a fast serial computer to simulate a parallel processing device.

^{xx} Such an account makes it in principle impossible for a differently embodied being to fully share human perceptual experiences. In this vein, the account has been accused (Clark and Toribio (2001)) of a kind of ‘sensorimotor chauvinism’. Noë (2004, p.26–28, and p. 113) rejects the charge of *unacceptable* chauvinism arguing that while every difference in sensorimotor contingencies must indeed make some slight difference to perceptual experience, there is plenty of room to allow that differently embodied beings may enjoy sufficiently similar experiences for each to count as, say, visual. It should be

noted that Noë can allow different surface features to yield the same experience (as when we adapt to inverting goggles: the 'bodily' surface is now altered but the experience normalizes after a while. What he does not allow (and see note 21 below) is that bodily differences that impact sensorimotor contingencies can be reconciled with full sameness of experience. Every such difference, for Noë, *makes* a difference.

^{xxi} Thus Noë (personal communication) does indeed assert that "you couldn't have the very same experience unless you have the same underlying sensorimotor exercise". This may turn out to be true, but it is not yet obvious to me why it *must* be true, or how we can at this time know it to be true.

^{xxii} For a review, see Clark (2003) chapter 4

^{xxiii} For a famous meditation on this theme, see Dennett (1981).