Pushing the Boundaries of Consciousness and Cognition

Kenneth Pepper

PhD

University of York

Philosophy

September 2014

<u>Abstract</u>

This thesis synthesises material from contemporary cognitive science, analytic philosophy of mind continental phenomenology to defend a view of the mind as embodied and extended. The first three chapters focus primarily on embodiment, while the last two chapters focus more on factors external to the body. In chapter I, I introduce Merleau-Ponty's concept of the body schema and argue that we should resist reducing the body schema to an internal representation of the body, and also that it does not always coincide with the boundaries of the biological body. In chapter II, I explicate and defend the sensorimotor approach to visual perception, further invoking Merleau-Ponty's phenomenology to support the arguments therein and to address certain worries internal to the sensorimotor approach. Chapter III builds on the conclusions of chapters I and II to explore one way in which technological extensions of the body can lead to novel perceptual experiences, and tentatively suggests a limited sense in which these experience may still be said to be visual in character. In chapter IV, I move beyond the body to explicate and defend the extended mind thesis, according to which cognition can and often does take place partly outside of brain and body via the active use of external aids and props. Finally, in chapter V, I consider the question of whether, given the perceptual phenomenology described in chapter II and the case for cognitive extension presented in chapter IV, consciousness might also be said to be extended, and argue that if certain assumptions are granted, it can.

Table of Contents

IV. Beyond the Body: Extending the Mind

V. Extended Consciousness?

Conclusions *237*

Acknowledgements

I would like to thank my supervisors Tom Baldwin and Keith Allen for all their invaluable advice, support and encouragement over the last four years, Tom Stoneham and Nick Jones for taking a risk by awarding me a graduate teaching scholarship, Clare Carlisle for convincing me (for better or worse) that it was all worth pursuing, and James Tartaglia and Josie D'Oro for inspiring me in the first place. I would also like to thank Mike Wheeler and Louise Richardson for agreeing to examine (and pass!) this thesis. Special thanks are also due to Andy Clark for providing me with a copy of an unpublished paper of his, and to various anonymous journal referees for their highly instructive comments. Additionally, I have benefitted enormously from the insightful discussions I have had over the years with Dom Shaw, Ema Sullivan-Bissett, Helen Bradley, Brendan Harrington, Owen Hulatt, Will McNeil and, especially, Simon Høffding. Finally, I would like to express my gratitude towards the third year students who elected to take my and Dom's Merleau-Ponty module, particularly Seb Odell, Ewan Margrave, Steve Startup and Susan Albinson, whose analytical and scholarly skills frequently put my own to shame, and helped shaped my own thoughts on many of the central issues of the thesis.

Completion of this thesis would not have been possible without the unwavering support of my partner Ally, my parents, my sister Sally, and my best friend Ian. I love you all for putting up with me for so long.

Author's Declaration

I hereby declare that the work herein is my own and that the work of others has been acknowledged appropriately.

Material from chapter II was published as 'The Phenomenology of Sensorimotor Understanding' in Mark Bishop and Andrew Martin eds. *Contemporary Sensorimotor Theory,* Dordrecht: Springer, 2014, pp.53-65.

Material from chapter V was published as 'Do Sensorimotor Dynamics Extend the Conscious Mind?' in *Adaptive Behaviour* vol. 22 no. 2, 2014, pp.99-108.

Draft material that would eventually form various parts of the thesis was presented at the British Postgraduate Philosophy Association's 15th annual conference at the University of Reading on the 24th of September 2011, the AISB Sensorimotor Theory workshop at Goldsmiths College London on the 26th of September 2012, the 11th meeting of the Nordic Society for Phenomenology at the University of Copenhagen on the 19th of April 2013, and the 1st Trends in Interdisciplinary Studies conference at Nicolas Copernicus University on the 7th of November 2013.

<u>Introduction</u>

1. Summary of the project

This thesis defends the claim that minds are embodied and, often, extended. Though we can and do engage in private reflection, the lion's share of our mental lives do not take place in our heads, but in the world, and this is owed to the way in which we are embodied and situated. To do so, it synthesises elements of analytical philosophy of mind, empirical and theoretical work in the cognitive sciences, and insights from continental phenomenology.

The claim that the mind is embodied and/or extended is not unique to this thesis. On the contrary, a good deal of it shall be occupied by the defence of arguments and claims already made by a variety of philosophers. Its intended original contribution is twofold. Firstly, each chapter contains counterarguments to objections to its central claims that, to the best of my knowledge, are not to be found in the extant literature. Secondly, it contributes to an emerging tradition that seeks to reconcile our lived experience with the findings of cognitive science, and illuminate the ways in which the two might be mutually informing and constraining (see e.g. Wheeler 2007, Gallagher 2005, Dreyfus 1992).

The structure of the thesis is as follows. I begin at the boundaries of the body, and gradually move outwards. The first three chapters focus primarily on embodiment, while the last two chapters focus more on factors external to the body. In chapter I, taking Husserl's account of the 'lived body' as my starting point, I introduce Merleau-Ponty's concept of the body schema and situate it in relation to empirical case studies of bodily disorders and adaptations. I argue that we should resist assimilating the

body schema to an internal representation of the body, and also that it does not always coincide with the boundaries of the biological body. In chapter II, I explicate and defend the sensorimotor approach to visual perception, further invoking Merleau-Ponty's phenomenology to support the arguments therein and to address certain worries internal to the sensorimotor approach. Chapter III builds on the conclusions of chapters I and II to explore one way in which technological extensions of the body can lead to novel perceptual experiences, and tentatively suggests a limited sense in which these experience may still be said to be visual in character. In chapter IV, I move beyond the body to explicate and defend the extended mind thesis, according to which cognition can and often does take place partly outside of brain and body via the active use of external aids and props. Finally, in chapter V, I consider the question of whether, given the perceptual phenomenology described in chapter II and the case for cognitive extension presented in chapter IV, consciousness might also be said to be extended. I argue that if certain assumptions are granted, it can.

2. Obligatory methodological preamble

This thesis is, by nature of its subject matter, a work in the philosophy of mind. However, throughout the thesis, particularly in the first three chapters, I will be appealing to arguments and descriptions found in literature from the European phenomenological tradition. I will make considerable use of Merleau-Ponty's *Phenomenology of Perception*, and shall also delve into the works of the discipline's founder, Husserl, as well as some limited appeal to Heidegger. My decision to appeal to these philosophers, as opposed to others who may share similar views, is largely a

pragmatic affair. For better or worse, these are the historical philosophers I have studied more than any others. As such they are the philosophers whose work has most shaped my own thinking on topics such as consciousness, perception and action. Rather than defend this in advance, it is hoped that the relevance of the aspects of phenomenological philosophy to which I appeal will be demonstrated by their utilisation in the substantive chapters that follow, but the essential guiding thought behind my appropriation of phenomenology is this: analytical philosophy of mind and the cognitive sciences face the task of accurately classifying and explaining the phenomena that purvey our mental lives. If these accounts rest on a phenomenologically inadequate conception of these phenomena, then they are flawed. Phenomenology is therefore an invaluable tool in guarding philosophy of mind against straw man conceptions of experience and behaviour by ensuring, in Husserl's words, that, insofar as this is possible, its accounts 'satisfy the principle of freedom from presuppositions' (Husserl 2001 p.97).

Though interest in phenomenology is steadily on the rise in contemporary Anglophone philosophy, it is still viewed with a mix of suspicion and outright disdain in some staunchly analytic circles. Though I suspect nothing I could say would suffice to dispel these anti-phenomenological sentiments altogether, the prevalence of such a negative bias does introduce an obligation to say something about what phenomenology is *not*. The use of the term 'phenomenology' as deployed in analytic philosophy generally refers to a certain kind of subject matter reserved for a particular domain of the philosophy of mind, namely, what is variably and interchangeably referred to as the 'qualia', 'raw feel', 'phenomenal character' or 'what-it-is-likeness' of conscious mental episodes. 'Phenomenology' in this sense designates a class of entities or properties of entities that it is the job of philosophers and

cognitive scientists to reconcile with our naturalistic worldview. I will on occasion also use the term in this way, and I hope that the context of usage is sufficiently instructive with respect to the term's desired sense. In the research program initiated by Husserl and continued by Heidegger, Sartre, Merleau-Ponty and others, however, 'phenomenology' refers not to a particular subject matter, but to a broadly shared approach to philosophical problems that takes as its starting point the first-person perspective of lived experience, prior to its being subjected to the abstract topographies inherited from particular scientific and/or philosophical theories.

Because of its emphasis on describing experience qua experience, there are philosophers who are wont to characterise phenomenology as a species of introspection. On such a conception, phenomenology should go the way of the introspectionist school of psychology associated with (inter alios) Brentano, Wundt and Titchener, long since superseded by behaviourism and now, thankfully, contemporary cognitive science. Hence, Dennett, who characterises Husserl's philosophy as advocating as 'a special technique of introspection' (Dennett 1991 p.44), dismisses the entire phenomenological movement thus: '[T]here really are no phenomenologists: uncontroversial experts on the nature of things that swim in the stream of consciousness' (ibid. pp.44-5). We can agree that there are no 'uncontroversial experts' in phenomenology, but this is true in the humanities and sciences generally. But Dennett's caricature of phenomenology as an introspective enterprise attacks a straw man, albeit one born of his quite justifiable worries about Husserl's insistence on philosophy's autonomy from natural science.[1] Husserl and subsequent phenomenologists reject this kind of Cartesian divide between an

---

[1] Deep down, I think Dennett knows this, and he does go some way toward acknowledging it in (Dennett 2007).

external world-in-itself and an internal realm of consciousness precisely because belief in such a divide is a), precisely the sort of ontological or metaphysical commitment to be deliberately bracketed out of phenomenological analysis and b), is not evident in experience; when I describe my perception of a tree, qua phenomenological description, it is an appearance of the tree itself that I am describing, and not a private mental image inside my head, cut off from the external world. The phenomenologist describes the world as it is experienced, not the inventory of some private stream of mind-stuff, as Husserl himself makes clear

> I perceive the physical thing, the Object belonging to nature, the tree there in
> the garden; that and nothing else is the actual Object of the perceptual
> "intention". A second immanental tree [i.e. an representation of a tree], or
> even an "internal image" of the tree, is in no way given, and to suppose that
> hypothetically leads to an absurdity. (Husserl 1999a p.90)

Indeed, the 'absurdity' to which Husserl refers in this passage is the very homunculus fallacy, and the infinite regress it incurs, that Dennett attacks in his critique of dualism.[2] Merleau-Ponty too, rejects introspection as a method for both phenomenology and psychology (Merleau-Ponty 1964a).

---

[2] 'The image as a really inherent component in the psychologically real perception would be again something real - something real that would function as a depicturing of another something real. But that can only be by virtue of a depicturing consciousness in which something first appears - with which we would have a first intentionality - and this would function again in consciousness as a "picture Object" representing another "picture Object" - for which a second intentionality founded in the first intentionality would be necessary [...] [If] perception and, then consequently, every mental process, requires a depictive function, [this] unavoidably (...) leads to an infinite regress. In contradistinction to such errors we have to abide by what is given in the pure mental process.' (Husserl 1999a. p.90-1)

Still, having dealt with these misgivings I must qualify my use of phenomenology somewhat. There is, of course, a huge array of deeply problematic aspects of individual phenomenologists' work. As with any philosophical school, differences and disagreements abound between its major figures (for a map of the conceptual terrain, see Moran 1999). I shall largely be ignoring the more problematic aspects of these phenomenologists' work and their multifarious points of disagreement in order to focus on their areas of overlap as they pertain to the topics of embodied and extended cognition. There is also a great wealth of technical detail in all of the major phenomenological texts I shall cite that are far beyond the scope of what might appear to scholars of phenomenology to be a rather superficial engagement, picking and choosing the relevant aspects of these philosophers' thinking to suit my own more philosophy of mind-oriented ends. My appropriation of phenomenology is therefore admittedly somewhat textually promiscuous, but it is undertaken towards the honest end of analysing mental phenomena with a degree of phenomenological accuracy often lacking in more orthodox analytical treatments. That phenomenology has much to offer the analyses of cognitive science is a belief I shall unapologetically leave undefended. Evidence enough for this can be found in the vast majority of articles published in the journal *Phenomenology and the Cognitive Sciences* since its inception in 2002, the fundamental ethos of which is beautifully illustrated and defended at length in Gallagher and Zahavi's (2007) seminal textbook, *The Phenomenological Mind*.

In addition to raiding the pantry of phenomenology I shall also, where appropriate, make extensive use of empirical work from the cognitive sciences, especially experimental psychology and cognitive neuroscience. Although I have no formal scientific training, I have endeavoured to comprehend the relevant empirical details

to the best of my ability and to deploy these in a way free from distortion, oversimplification or overstatement. The more empirically-oriented parts of the thesis are largely necessitated by the data-driven nature of much of the philosophical literature I shall be engaging with. But also, just as I cannot bring myself to take seriously any philosophical or psychological analysis which rests on a phenomenologically false account of the phenomena under investigation, I find it virtually impossible to see how a satisfactory account of any mental phenomenon can be genuinely achieved without due care and attention to at least some of the fascinating and deeply pertinent insights to be garnered from the rapidly advancing field(s) of contemporary cognitive science. A philosopher who pursued such a line of research would, in my view, be in the same boat as an art historian who avoided viewing art or a literary critic who neglected to read literature.

Phenomenology is not at odds with cognitive science. As Merleau-Ponty himself remarks 'Psychology and philosophy are nourished by the same phenomena; it is only that the problems become more formalized at the philosophical level' (Merleau-Ponty 1964b p.24). Though phenomenological analysis might *consequently* prompt us to question some of the more philosophical claims made by scientists or philosophers on their behalf, if say, in the light of phenomenological considerations these claims seem to presuppose a phenomenologically false characterisation of the phenomena they purport to explain, we need not and should not doubt or deny any scientifically-established facts in order to describe experience qua experience, just as we should not contradict ordinary experience for the sake of rendering an account or explanation more palatable to naturalistic dogma. Hopefully in what follows I shall do neither.

I. Being-in-the-World and the Body Schema

1. Introduction

This chapter introduces the general theme of embodiment. I begin with a brief discussion of Husserl's account of the lived body and its shortcomings, paving the way for a phenomenological outlook that is closer in spirit to that of Heidegger and Merleau-Ponty. I then introduce the concept of the body schema, which will play an important role in chapters II and III. This is subsequently contrasted with the related but different notion of a body image, and I distinguish between two notions of 'representation' as it pertains to this distinction. I then discuss some insights from cognitive neuroscience to support the claim that the body schema, as I have defined it, is flexible with respect to the boundaries of the organism. I argue that we should resist assimilating the body schema to underlying neural states or processes.

2. Making room for the body

Traditionally, philosophy has struggled to know quite what to do with the body. In Plato's *Phaedo,* Socrates tells Simmias that for true philosophers, the body is a distracting obstacle standing in the way of the pursuit of knowledge and truth, even going so far as to suggest that only upon death, defined as the separation of the psyche from the body, does genuine philosophical insight become possible. In the meantime,

> While we live, we shall be closest to knowledge if we refrain as much as possible
> from association with the body and do not join with it more than we must, if we

are not infected with its nature but purify ourselves from it until the god himself frees us. (Plato 2002 p.104)

Fortunately for we mere mortals, the asceticism Socrates advocates with respect to the body is a barely conceivable philosophical fantasy. Beyond the basic biological necessities of digestion and respiration, the mental lives of human subjects - their perceptions, actions, and experiences - are profoundly and inescapably intermingled with, and highly dependent upon, their embodiment. But this statement, perhaps, is uninformative and uncontroversial for even the most ardent of mind-body dualists. Even Descartes, for whom minds and organisms are two entirely distinct substances, happily admits that 'it is not enough for [the mind] to be lodged in the human body like a pilot in his ship' (Descartes 1998 p.33). To fully appreciate the scope of the claim that we are embodied, we need to illuminate what it means to be a bodily subject, rather than a subject that is merely housed in a body. In essence, this is the project of Merleau-Ponty's landmark text *Phenomenology of Perception*.

Husserl's discussion of the body in the second volume of his *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy* is an important precursor to the phenomenology of embodiment developed by Merleau-Ponty. Though I shall not attempt to explicate it fully, a few key points of Husserl's understanding of embodiment provide a natural starting point for this chapter. For Husserl, the Body is 'doubly constituted', meaning that we can experience our own bodies in either of two ways. One way in which we experience the body is as an object of external perception bound up in the causal nexus of other material things. Another way in which we experience the body is as the field of sensitivity where bodily sensations occur, and as a unique object which we can move at will from the inside. Hence, Husserl distinguishes between the *lived* body (Leib) - the animate

body as experienced from the inside, and the *objective* body (Korper) - the body understood as a mere material object among others. Husserl's distinction is echoed by Merleau-Ponty's own distinction between 'phenomenal' and 'objective' conceptions of the body[3].

Husserl's interest in the body is motivated by the question of how perceiving subjects can experience successive percepts as presentations of one and the same object (Zahavi 2003 p.99). Husserl's answer is that successive visual presentations of objects are endowed with their unified appearance, which he calls their 'schemata', by virtue of their occurring simultaneously with schemas of *kinaesthetic* sensations localised in the body

> Those sensations which undergo extensional apprehension (leading to the extended features of the thing) are motivated as regards the courses they take either actually or possibly and are kinaesthetically *related to motivating series, to systems, of kinaesthetic sensations,* which freely unfold in the nexus of their familiar order in such a way that if a free unfolding of one series of this system occurs (e.g. any movement of the eyes or fingers), then from the interwoven manifold as motive, the corresponding series must unfold as motivated…We constantly find here this two-fold articulation: kinaesthetic sensations on the one side, the motivating; and the sensations of features on the other, the motivated. (Husserl 1989 p.62-3, italics in original)

---

[3] 'We never move our objective body, we move our phenomenal body, and we do so without mystery, since it is our body as a power of various regions of the world that already rises up toward the objects to be grasped and perceives them.' (Merleau-Ponty 2012 p.108)

Husserl's understanding of 'kinaesthesia' is somewhat broader than contemporary technical usage. As well as the sense of bodily movement, Husserl seems to have in mind what we would now refer to as proprioception (the sense of bodily positioning and posture) and equilibrioception (the sense of balance, also referred to as the vestibular sense). Husserl argue that every presentation of an object's perceptible features (colour, surface, aspect, etc) - a 'motivated sensation' - must be accompanied by a 'motivating sensation' localised in or on the lived body, and it is only by virtue of the localisation of sensations that a subject can have a lived body, because it is by virtue of the localisation of bodily sensations that the body imposes a perspective on the perceived world

> Obviously, the Body is also to be seen just like any other thing, but it becomes a *Body* [Leib] only be incorporating tactile sensations, pain sensations, etc - in short, by the localisation of the sensations as sensations. In that case the visual Body also participates in the localisation, because it coincides with the tactual Body, just as other things (...) coincide, ones which are constituted both visually and tactually, and thus there arises the idea of a sensing thing which "has" and which can have, under certain circumstances, certain sensations (...) and, in particular, have them as localised in itself primarily and properly. (Husserl 1989 p.158-9)

The phenomenological significance of the lived body, for Husserl, then, is that it provides the 'zero point' of all perception by virtue of its being the bearer of localised tactile and/or kinaesthetic sensations (ibd. p.166). We can therefore discriminate two distinct claims in Husserl's treatment of the body. The first is that all perceptual contents are accompanied by kinaesthetic sensations, broadly construed. The second is that a body is only a *lived* body insofar as it bears localised sensations. The first

claim is false. There are plenty of people who enjoy perfectly normal perceptual experience who are physically incapable of experiencing anything corresponding to Husserl's notion of kinaesthetic sensation, even if this is read liberally to include proprioception and equilibrioception. Paralysis is an obvious example, and I shall discuss the extreme case of locked in syndrome in the next chapter to argue that their lack of bodily *sensation* does not suffice to discredit the more general claim that perception is fundamentally an embodied phenomenon. The second claim is of more immediate interest in the context of this chapter.

There is far more to being embodied than feeling sensations. Husserl's emphasis on sensation leads him to neglect the ways in which our experience of the world is structured *pre-noetically,* that is, prior to explicit conscious awareness, by the organisation of the body (Gallagher 2005 p.24). Hence, according to what I take to be a fairly standard reading, though a step in the right direction, Husserl's account of the lived body suffers from an overemphasis on the ways in which we can attend to our bodies and an overstatement of the extent to which the body shows up as an object of immediate awareness in our mental lives, due his overly mentalistic methodology (Dreyfus 2000; 1992, Carman 1999, cf. Smith 2007; Zahavi 2003 pp.98-109). Husserl's analyses takes place within his phenomenological reduction, the outcome of a systematic bracketing of the empirical and ontological commitments of our commonsense beliefs about the natural world which is undertaken in order to describe experience directly, that is to describe experience *qua* experience, which he explains thus:

> This universal depriving of acceptance, this "inhibiting" or "putting out of play" of all positions taken toward the already-given objective world and, in the first place, all existential positions (those concerning being, illusion,

possible being, being likely, probable, etc), -or as it is also called, this "phenomenological epoché" and "parenthesising" of the Objective world- therefore does not leaves us confronting nothing. On the contrary we gain possession of something by it; and what we (or, to speak more precisely, what I, the one who is meditating) acquire by it is pure living, with all the subjective processes making this up, and everything meant in them, *purely as meant* in them: the universe of "phenomena" in the particular (…) phenomenological sense. (Husserl 1988 p.20-1, italics in original)

Perhaps as a result of this self-professed Cartesian method (Husserl 1988), Husserl's analyses are largely restricted to the contents of particular mental acts. By reducing experience to the 'purely meant' intentional objects and their corresponding cognitive acts, Husserl in effect removes the subject under investigation from the practical context in which their behaviour occurs and in which they exercise a great deal of situational understanding. In essence, this is the implicit critique of Husserl's whole project that runs throughout Heidegger's *Being and Time*. Heidegger draws attention to the fact that Husserl's Cartesian method of systematic belief suspension, which even in his analyses of social phenomena and the development of modern science always leads back to a constituting transcendental ego (see, e.g. ibid.89-150), is impotent to accommodate the more basic, pre-reflective ('primordial' in Heidegger's terminology) level of experience human subjects enjoy prior to engaging in the sort of reflective theorising which enables metaphysics, natural science, and the would-be 'rigorous science' of Husserlian phenomenology (Husserl 1970). This underlying 'being-in-the-world', as Heidegger calls it, is the true default starting point for any genuine account of lived experience. In other words, phenomenology should not bracket away the practical world in order to uncover the conscious subject behind it,

but illuminate the nature of this subject by describing and exploring the ways in which it is inescapably embodied and embedded in its world

> When Dasein [i.e. beings like us, literally 'being-there'] directs itself towards something and grasps it, it does not somehow first get out of an inner sphere in which it has been proximally encapsulated, but its primary kind of Being is such that it is always 'outside' alongside entities which it encounters and which belong to a world already discovered. Nor is any inner sphere abandoned when Dasein dwells among the entity to be known...but even in this 'Being-outside' alongside the object, Dasein is still 'inside', if we understand this in its correct sense; that is to say, it is itself 'inside' as a Being-in-the-world which knows. (Heidegger 1962 p.89)

The most basic mode of encounter human beings have with their environment is a pre-reflective familiarity rather than an interrogative attitude. Whenever we begin to withdraw from our practical engagements with the world by reflecting on our experience and behaviour, we find ourselves always-already in a situation of sorts, embedded in a network of familiar practices which is lived through prior to being made thematic as an object of analysis. When we do engage in the activity of interpreting our experience of the world and our behaviour in it, we do not, as Heidegger puts it, 'throw a signification over some naked thing' by 'sticking a value on it' (ibid. p.190), but rather try to make explicit the tacit practical understanding that was already present in our dealings with our surroundings (ibid. p.191). Hence,

in what could just as well be a slogan for the research program of embodied cognition[4], Heidegger summarises the point thus

> The kind of dealing which is closest to us is...not a bare perceptual cognition, but rather that kind of concern which manipulates things and puts them to use; and this has its own kind of knowledge. (ibid. p95)

This Heideggerian sentiment - that phenomenology cannot succeed by abstracting the subject out of their world, is taken up by Merleau-Ponty and is the meaning of his remark in the preface to *Phenomenology of Perception* that '[t]he most important lesson of the reduction is the impossibility of a complete reduction' (Merleau-Ponty 2012 p.Ixxii). Hence, in a review of Husserl and Merleau-Ponty's respective accounts of the lived body, Carman writes

> In his posthumous works Husserl calls attention to the role of the body in perception, but he takes it for granted that cognitive attitudes rather than bodily skills must bridge the intentional gap between mind and world. He therefore attempts to ground bodily self awareness in what he takes to be a more basic form of intentionality: the quasi-objective localization of subjective tactile sensations in the body. But to tie the body's intentional constitution specifically to the sense of touch in this way [...] amounts to a fundamental misunderstanding of its significance for phenomenology. (Carman 1999 p.206).

---

[4] Strikingly, Heidegger himself has virtually nothing to say about the body, with the exception of the writings published as *The Zollikon Seminars,* where he applies the project of *Being and Time* in relation to certain neuropathologies (Heidegger 2001).

For present purposes, it doesn't really matter whether the Carmanean interpretation of Husserl I have espoused is accurate or, if not, what Husserl's precise position might have amounted to. The moral is that the errors of such a view of the body are instructive with respect to an adequate account of embodiment. To get a better grip on the ways in which our ordinary experience is structured by our embodiment, we need to make a concerted effort to avoid treating the body as though it were either a continual object of perception, or as something an explicit awareness of which accompanies all our experience. Of course, the body can be an object of a mental state. We can and do visually and tactually inspect our bodies, we can think about it in myriad ways, and it routinely shows up in our experience when we are injured or trying to perfect a new skill. But, in the ordinary flow of things, the body, qua intentional object, is highly recessive, and what bodily awareness there is when immersed in some behaviour or other is generally non-observational in character (Marcel 2003). As Wittgenstein famously remarked, 'voluntary movement is marked by the *absence* of surprise' (Wittgenstein 1958 p.162, my italics). They key to doing phenomenological justice to our experience, without skipping over the essential contributions of the body, then, is to articulate the ways in which embodiment structures experience pre-reflectively. In the next section, I introduce the concept of the body schema, which forms an integral part of such an account. The body schema will feature prominently in the discussions of perceptual experience and the senses that follow in chapters II and III.

3.   **The scheming body**

Prior to engaging in reflection on experience, the human (or animal) organism is always-already attuned to its environment. Our encounters with the world are made possible by having at our immediate disposal certain motor habits which alleviate the requirement for acts of judgement on our part. Ceteris paribus, people act without needing to consciously think through or attentively observe each individual limb movement. The individual bodily movements implicated in the performance of simple actions like grasping a cup to take a drink or taking a phone from one's pocket to answer it occur without attention or perceptual monitoring. Insofar as we are conscious of our actions, our experience is one of general bodily movement and a directedness towards the wider project at hand ('drinking', "answering the phone', etc), but not as a linear succession of sub-routines in accordance with a mentally represented list of instructions.

Dreyfus (1992) calls this kind of immersed bodily engagement with the world 'absorbed coping'[5]. Neurotypical[6] subjects engaged in absorbed coping will find that their hand and fingers naturally shape themselves around the object to be grasped without the need to feel their way around it. This is because they have at their disposal a system of either learned or innate motor abilities governed by low-level proprioceptive, kinaesthetic and vestibular processing that renders explicit judgements concerning their hand and arm movement redundant. They merely act on their impulse to grasp the coffee cup and before they know it (so to speak), their

---

[5] There is a whole host of problems with Dreyfus's developed account of the phenomenon of absorbed coping (see e.g. Hoffding forthcoming, Romdenh-Romluc 2007), but the term itself is helpful, so I shall continue to use it with the proviso that it is Merleau-Ponty's phenomenology, and not Dreyfus's, that I have in mind.

[6] The term 'neurotypical' was originally coined by the autistic community to refer to non-autistics, but now enjoys a broader usage to denote anyone fortunate enough to have a normally functioning brain. I thank Susan Albinson for teaching me this term.

hand is being raised accurately towards their mouth. As Merleau-Ponty stresses, the organisation of the body is such that *it,* and not the subject qua detached observer, orients itself in relation to the object to be grasped before any reflective perceptual judgement is made

> [M]y movements anticipate directly their final position, my intention only sketches out a trajectory in order to meet up with goal that is already given in its location, and there is something like a seed of movement that only grows later though its objective trajectory [...] I have no need of looking for it because it is always with me. I have no need of directing it toward the goal of the movement, in a sense it touches the goal from the very begging and throws itself towards it. (Merleau-Ponty 2012 p.96-7)

Merleau-Ponty calls this bodily system by which this happens the *body schema.*[7] The origin of the term 'body schema' is commonly attributed to classical neurologists Henry Head and Gordon Holmes, who coined it to denote a continually-updated internal map of bodily posture

> Every recognizable change enters into consciousness already charged with its relation to something that has gone before [...] So the final product of the tests for the appreciation of posture or passive movement rises into consciousness as a measured postural change. For this combined standard, against which all subsequent changes of posture are measured before they enter consciousness, we propose the word "schema". By means of perpetual alterations in position we are always building up a postural model of ourselves which constantly

---

[7] '[M]y entire body is not for me an assemblage of organs juxtaposed in space. I hold my body as an indivisible possession and I know the position of each of my limbs through a *body schema* [*un schema corporel*].' (ibid. pp.100-1, translator's italics)

changes. Every new posture or movement is recorded on this plastic schema, and the activity of the cortex brings every fresh group of sensations evoked by altered posture into relation with it. (Head & Holmes 1911 p.187)

The term is adopted and modified by Merleau-Ponty to denote not an internal mental map of the body, but an integrated network of motor capacities that structure perception and action. In Merleau-Ponty's usage, the functioning of the body schema constitutes a fundamental kind of practical understanding that is exercised by subjects pre-reflectively. The body schema is an enabling condition of absorbed coping; without prior deliberation, the body comports itself toward the world.

By denying that absorbed coping is the outcome of a judgement, or a series of judgements, we do not thereby lapse back into a behaviouristic view of action, whereby the bodily movements required to (say) grasp the cup on the basis of visual perception are mere stimulus-response reflexes. (Merleau-Ponty calls these views of action 'intellectualist' and 'empiricist' respectively.) Watson, for example, famously held that psychology 'can be done in terms of stimulus and response, in terms of habit formation, habit integrations and the like' alone (Watson 1913 p.167). What the behaviourist misses, according to Merleau-Ponty, is the fact that reflexive movements of this kind 'are never blind processes: they adjust to the "sense" of the situation' (Merleau-Ponty 2012 p.81). Bodily responses to our surroundings are already meaningful for the agent; they make practical sense given the demands of one's present situation as one perceives it. We instinctively raise our hands to cover our face when an object is thrown in our direction, brace ourselves when we fall, turn our heads upon hearing a crashing sound behind us, and so forth. Hence, channelling

Heidegger, Merleau-Ponty writes, 'Reflex, insofar as it opens itself to the sense of a situation, and perception, insofar as it is an intention of our total being, are modalities of a *pre-objective perspective* that we call "being in the world' (ibid.). On Merleau-Ponty's account it is therefore habit, not thought, which is primarily constitutive of our being-in-the-world

> [I]n the acquisition of a habit, it is the body that "understands". This formula will seem absurd if "understanding" is the act of subsuming a sensory given under an idea, and if the body is a mere object. But the phenomenon of habit in fact leads us to rework our notion of "understanding" and our notion of the body. To understand is to experience [...] the accord between what we aim at and what is given, between the intention and the realisation - and the body is our anchorage in the world. (Merleau-Ponty 2012 pp.145-6).

This implicit 'sense' of motor habits is indicative of a primordial grasp on the world. Our body schematic routines orient themselves in accordance with the motor tasks in which we engage - making a cup of coffee, answering a buzzing phone, defending ourselves from blows, playing the guitar, etc. The body schema therefore exhibits its own kind of 'world-directness' or 'aboutness' in the absence of a pre-conceptualised goal, that is, its own particular kind of intentionality, which Merleau-Ponty calls *operative* or *motor intentionality* (ibid. p.112-4). As I read Merleau-Ponty, operative intentionality consists in a minimalistic intentional relation between subject and world which does not have a determinate object, but is dependent for its content on a present context of action, and operates largely below the level of conscious attention or explicit awareness. This contrasts with what might be called the 'object' intentionality which Husserl (e.g. 2001) following Brentano, emphasises as the defining feature of consciousness. This kind of intentionality involves having a

determinate object present to the mind via a mediating cognitive act. Here we recall the oft-quoted 'Brentano's thesis'

> Every mental phenomenon is characterized by what the Scholastics of the Middle Ages called the intentional (or mental) inexistence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction toward an object (...),or immanent objectivity. Every mental phenomenon includes something as object within itself, although they do not all do so in the same way. In presentation something is presented, in judgement something is affirmed or denied, in love loved, in hate hated, in desire desired and so on. (Brentano 1995 p.68)

Object intentionality is the kind of intentionality proper to propositional thought. Operative intentionality, on the other hand, provides the intelligible background against which our own behaviour unfolds prior to thought. It can therefore be glossed as a kind of non-conceptual or pre-conceptual intentionality (Jensen 2009). For reasons clarified below, however, unlike some philosophers I shall resist describing it as non-representational (see e.g. Mooney 2011, Romdenh-Romluc 2011a, Jensen 2009). There are many philosophers for whom 'representation' and 'intentionality' are synonymous (e.g. Crane 2009; 1995), so the notion of 'non-representational intentionality' borders on the oxymoronic and has the potential to create unnecessary confusion. Moreover, although the operative intentionality of the body schema is certainly non-representational in the sense in that, unlike object intentionality, it does not represent any particular object or fact, there are other, more empirical, technical uses of 'representation' which are germane to the concept of operative intentionality, which I discuss below. First though, I juxtapose the notion

of a body schema to that of a body image, which serves to further illustrate the distinction between object and operative intentionality as it pertains to the body.

4. Body schema vs. body image

The body schema, as I am using the term, differs from what is popularly called a 'body image'. As Gallagher notes, the terms 'body schema' and 'body image' are often used interchangeably in discussions of the body across academic disciplines (Gallagher 2005 ch. I; 1986a). This is unfortunate, as there are sound phenomenological, conceptual and empirical grounds for distinguishing between the two (ibid.). To avoid conflation, Gallagher provides useful criteria by which to make this distinction,

> A *body image* consists of a system of perceptions, attitudes and beliefs pertaining to one's own body. In contrast, a *body schema* is a system of sensory-motor capacities that function without awareness or the necessity of perceptual monitoring [...] A body image [...] can include mental representations, beliefs, and attitudes where the object of such intentional states (...) is or concerns ones' own body. The body schema, in contrast, involves certain motor capacities, abilities and habits that both enable and constrain movement and the maintenance of posture. It continues to operate, and in many cases operates best, when the intentional object of perception is something other than one's own body. (Gallagher 2005 p.24, italics in original)

Following Gallagher, then, we can make the following summarisation. To have a body image is to have a mental state or perception, the object of which is one's own

body. For example, forms of body dysmorphia, a form of anxiety disorder, are essentially disorders relating to a subject's body image. People suffering from anorexia are plagued by the worry that they are overweight (for a clinical review, see Morris & Twaddle (2007)), while people suffering from muscle dysmorphia (colloquially known as 'bigorexia'), which is especially prevalent among male bodybuilders, suffer from anxiety over a perceived lack of muscle mass (see e.g. Mosley 2009). The anxieties of body dysmorphics stem from how they perceive or feel about their bodies and/or how they judge them to be relative to an idealistic aesthetic standard, that is, how their bodies are given to them as objects of thought or perception. The body schema, is contrast, is neither a mental state nor the content of a mental state, but a network of largely unconscious integrated bodily capacities.

Reflection on body dysmorphia bolsters the conceptual distinction between body image and body schema. The body dysmporphic bodybuilder has misconceptions about the size and shape of his body, but while slaving away in the gym has no trouble coordinating his actions. On the contrary, through his training he builds up bodily habits which enable him to perfect the correct form of the exercise(s). His false beliefs and erroneous perceptual judgements about his own body do not prevent him from, say, accurately grasping a dumbbell and expertly performing repetitions of an exercise. And anorexics, who cannot shake the feeling that their bodies are too large and are prone to misjudge their size and build when, say, looking in a mirror, do not overcompensate for their perceived additional body mass in their daily negotiations with their surroundings. Despite their disorder of body image, their body schemas remain intact and unaffected.

The case of Ian Waterman, who lost touch and proprioception below the neck to acute sensory neuropathy caused by a rare infection, provides further, empirical,

support for the schema/image distinction.. His case has been documented at length

by clinicians, psychologists and philosophers (Gallagher 2005, Cole, Gallagher &

McNeil 2002, Cole 1995, Cole & Paillard 1995, Gallagher & Cole 1995) and even a

BBC documentary, 'The Man Who Lost His Body', which at the time of writing is

still available on Youtube[8]. Upon his initial hospitalisation, Waterman could not

even sit upright. If he tried to move one of his limbs it would flail wildly and the rest

of his body would also move in unpredictable ways. Though his muscles still worked

and he could tense them, he had lost the prioprioceptive feedback that informs his

brain where his body parts are. Eventually, through an amazing feat of willpower,

Waterman learned to not only sit upright, but to live an active life that to casual

observers, appears close to normal. He walks, drives, works and pursues hobbies. But

his nerves remain severed. He succeeds in controlling his movement by the use of an

often exhaustingly high degree of visual attention, and maintains posture by tensing

his muscles and relying on memory of where he last saw his limbs to be (Gallagher

2005 pp.43-5). Though Waterman gestures seemingly normally in conversation,

experiments reveal that this too relies on close visual monitoring. Cole, Gallagher

and McNeil (2002) asked Waterman to watch a short cartoon and relay the narrative.

When describing the events depicted in the video he gesticulated with his hands in an

apparently ordinary manner. But when set the same task with his hands shielded

from view he did not, the absence of visual feedback preventing him from doing so,

and kept his hands clasped in his lap (ibid.). Waterman relies on his preserved

capacity to feel sensations of heat and uses these to clasp his hands together to keep

them still (Gallagher 2005 p.112).

---

[8] https://www.youtube.com/watch?v=AvLRwRAAoww (last accessed 31/8/14).

Waterman's deficiency consists not in an inability to make judgements about his bodily location and positioning as he can do so on the basis of observation and inference. Rather, he has lost the low-level bodily awareness that typically renders such judgements unnecessary, owing to an almost total loss of body schema. (I say 'almost', because Waterman has retained proprioception in his head and face, and is still able to tense his muscles, allowing him to freeze his position when, for example, sitting upright. He has also retained some visual proprioception - feedback about bodily movement and posture via visual information). In its place, he relies heavily on a continually maintained (perceptual) body image, as Gallagher explains

> In place of the missing schema processes, we might say that Ian has substituted a virtual body schema - a set of cognitively driven motor processes. This virtual schema seems to function only within the framework of a body image that is consciously and continually maintained. If he is denied access to a visual awareness of his body's position in the [visual] perceptual field, or denied the ability to think about his body, then, without the framework of the body image, the virtual body schema ceases to function - it cannot stand on its own. (Gallagher 2005. p.52)

Waterman's body must always feature in his experience as a third term between his intention to act and the world in which he acts. It wouldn't be strictly correct to claim that his body is now entirely devoid of the operative intentionality Merleau-Ponty ascribes to the body schema, as he does not need to consciously think through individual muscle contractions, but he does have to attend closely to what each body part is doing in order to judge whether or not his movements correspond to his intentions (ibid. p.52-3) and he has lost the ability to form new motor programs (ibid.

p.62). As such Waterman's body is only present for him when it is the intentional object of an occurrent conscious mental state (ibid. p.64).

Sacks (1985) describes an even more extreme case. A patient of his referred to as 'Christina' suffered a complete loss of proprioception caused by sensory neuritis on the roots of her spinal and cranial nerves. Like Waterman, Christina can now only locate her limbs by continually visually attending to them. In contrast to Ian Waterman, her loss of proprioception extended to her head, making speech modulation initially impossible, though she eventually learned to reply on auditory feedback. Through the use of strict visual attention to her movement and posture, Christina regained the ability to coordinate her body, return to work, and lead what can appear to outsiders to be a relatively normal life, though her movements and speech now look like unnatural 'theatrical performances' (ibid. p. 54) due to the exaggerated nature of her visually-guided compensatory movements. On hearing her prognosis, Christina remarks

> I've already noticed…that I may 'lose' my arms. I think they're one place, and I find they're another. This 'proprioception' is like the eyes of the body, the way the body sees itself. And if it's gone, as it's gone with me, *it's like the body's blind.* (ibid. pp.51-2)

Again, a body percept must be continually maintained in order to compensate (indeed, overcompensate) for the loss of body-schematic motor routines. Christina goes so far as to describe herself as 'disembodied' but because she no longer identifies with her body as hers, and struggles to remember or even imagine (ibid. p.56), indicating one way in which a lost body schema can have a profound long-term effect on a body image. In a dramatic choice of words, Sartre compares the

objective conception of the body as a body modelled on the anatomical understanding of a corpse (Sartre 2003 p.327). Christina, who describes herself as 'pithed like a grog' (Sacks 1985 p.56) presents one case where the phenomenal body has for her, in a certain sense, died. Both cases illustrate the extent to which the body schema, though generally recessive, plays a vital structural role in our mental lives.

5. 'Representation' and 'representation'

Neuroscientists and cognitive psychologists, influenced by Head and Holmes, generally use the term 'body schema' to refer to a neurally encoded internal map of the body (see e.g. Cardinali et al 2009, Marivita & Iriki 2004, Maravita, Spence & Driver 2003). Here one recalls the famous model of Penfield's 'sensory homunculus', an anthropomorphic visual illustration of the human body in proportion to the areas of the somatosensory cortex dedicated to each body part (Penfield & Rasmussen 1950). The hands, genitals, and lower part of the head (especially the lips) of the model are disproportionately (and amusingly) larger than a real human body owing to the greater cortical space dedicated to their sensitivity (Ramachandran & Blakeslee 2005 p.25-7). This inner 'body map' is what most neuroscientists have in mind when they talk about a body schema - an internal map of the body that the brain uses to control bodily movement.

Following suit, de Vignemont (2010) challenges Gallagher's characterisation of the case of Ian Waterman as one of a 'missing schema':

> It does not make more sense to claim that deafferented patients have a deficit of body schema than to claim that blind people have a deficit of body schema. In the latter case, they rely on proprioception instead of vision. In the former case, they

rely on vision instead of proprioception. The body schema (qua sensorimotor representation) is there, although based on different weighting of information. Whereas proprioception may normally play an important role, it has been taken over by vision in deafferented patients. One may reply that their actions are not normal because they require reflexive monitoring of their movements. This might have been true at the beginning when the patients had to learn how to visually guide their movements. Similarly, you have to pay attention at the beginning when you learn to drive a car. But after a while, you drive without consciously monitoring the visual information that you receive on the other cars, the road, etc. Deafferented patients are like the automatic drivers. [...] This is not to say that their body schema has not changed. It has, giving more weight to vision. But it is not "missing". (de Vignemont 2010 p.675)

It is a fair but trivial point that 'The Case of the Missing Schema' is an ever-so-slightly inaccurate title for the second chapter of *How the Body Shapes the Mind*, as in it Gallagher actually argues, as aforementioned, that a), Waterman suffers from an *almost* total, but not *totally* total, loss of body schema, as he has retained proprioception and tactile sensation in his head as well as visual proprioception and is still able to freeze his bodily position at will, and b), by learning to control his posture and movement via visual attention, he has developed a 'virtual' body schema. But this concession, if it even is a concession, does not license the claim that a patient such as Waterman has a body schema in the sense in a blind but otherwise able bodied subject can be said to have one. De Vignemont reaches this conclusion by misrepresenting Gallagher's position by identifying the body schema with a 'sensorimotor representation'. Recall that this is not the neo-Merleau-Pontian concept

of body schema with which Gallagher operates. Though the body schema is constituted, in part, by. sensorimotor representations, it is not itself a representation but a network of bodily dispositions, and it is precisely Waterman's loss of these motor dispositions, which in neurotypical subjects alleviate the need for visual monitoring, that explains his behavioural and phenomenological atypicality. For Waterman his body must be, and can only be, the object of a percept or mental state. What matters here is that the operative intentionality of the body - its pre-reflective world-directedness, is vastly diminished. It is not the loss of proprioception per se which justifies the claim of a loss of schema, but the loss of motor dispositions that occurs as the direct result of the loss of proprioception; the type of information processing involved is a secondary concern. And the fact (assuming it is a fact) that Waterman has retained a neural representation(s) despite his loss of proprioception only serves to demonstrate that the body schema proper is not identical or reducible to a neural state, but is constituted by processes in the peripheral nervous system also as it is the connection between the two that is damaged (Gallagher 2005 p.48).

De Vignemont also errs by conflating the (true) claim that Waterman does not need to continually relearn how to use visual information to control his movement with the (false) claim that he does not consciously and reflexively visually monitor his movement. He does, he just does it very well. While it is true that his movements can seem automatic or close to automatic from outside observation, as a matter of fact they are not, and unlike those of de Vignemont's seasoned driver, require a high degree of conscious monitoring. Even for the well established activity of walking on a flat surface, Waterman estimates that it requires fifty to seventy percent of his attention, while walking on an uneven surface requires a hundred percent (Gallagher 2005 p.49). Waterman's need for a high degree of conscious visual attention to his

body is further evidenced by the fact that if he is asked to hold an egg he can do it apparently quite easily, but once his attention is directed elsewhere, he crushes it (ibid. p.44). Moreover, de Vignemont's comparison between Waterman's control of his body and driving a car is discredited by Waterman's own experience of driving; it is precisely the relative lack of conscious effort required to keep a car stable and moving that enables him to drive far more easily than walk

> Ian drives and he enjoys it. It seems effortless to him in comparison to walking. He reports that it is easier for him to drive 300 or 400 miles than to stop and refuel. Driving actually allows him to relax his attention to his bodily movements. Ian maintains his posture by 'freezing' in place. He needs only to keep his hands within the visual field (all the controls are manual), and he's assisted in a high degree of visual perception which facilitates automatic control in the case of driving (...). (ibid. p.58)

Waterman's car won't change shape or fall over if he looks away, and it is mechanically disposed to respond appropriately to its controls. Ian thereby offloads much of his mobility issues onto his 'car schema', as it were. Far from providing an analogy by which to ascribe a body schema to deafferented subjects like Waterman and Christina, de Vignemont's driving example only serves to further illuminate the stark differences between ordinary action and their highly vision-dependent coping strategy.

It is perhaps the guiding metaphor of an image of the body encoded in the brain - a sensory homunculus - which motivates theorists of various stripes to posit accompanying bodily mental imagery in their explanations of our experience and behaviour (e.g. Metzinger 2009; Damasio 2000, O'Shaughnessy 1980), and which

leads neuroscientists to assimilate the body schema to a neural representation. The assumption would be that as the brain processes representations of the body and/or various body parts, our ability to act skilfully in our surroundings must depend on a mental image of our body. This inference is unjustified, as it trades on an ambiguity between two senses of 'representation'. Unlike certain philosophers, I do not think that the notion of representation and its liberality of usage amounts to a call for its outright abandonment (cf. Hutto & Myin 2013, Hutto 2013); when deployed carefully it is a perfectly innocuous and very useful term. It does, however, necessitate some clarification as to the precise role of representation in the concept of the body schema I am defending here, and will appeal to at various junctures throughout this thesis.

It is prima facie tempting to read both Merleau-Ponty and Gallagher as rejecting all forms of representational talk in relation to the body schema, so that it might be distinguished from a body image, but this is not strictly warranted once their claims are put into context. We can distinguish between what we may, following Dennett (1969), call 'personal' and 'sub-personal' senses of 'representation'. The personal level of description pertains to what Sellars (1997) calls the 'manifest image' of humankind, or what Husserl calls the 'life-world' [Lebenswelt], which comprise the humanistic domain of our everyday discourse about our mental lives. The personal level of description comprises both folk psychology and the personal level analysis par excellence, phenomenology. In the personal level sense of 'representation', to represent something - a thing, a state of affairs - is to be in a mental state which has that thing or state of affairs (or whatever) as its object (a la object intentionality). Explanations of behaviour at the personal level appeal to people's beliefs, desires, intentions, perceptions and conscious experiences. In contrast, discourse at a sub-

personal level refers to what Sellars calls the 'scientific image' and which belongs to the domain of what Merleau-Ponty derisively calls 'objective thought'. Explanations of mental phenomena at a sub-personal level of description require a switch of theoretical stance from everyday discourse on experience in favour of a piecemeal analysis of the physical states and processes that underpin personal-level mental phenomena - neurobiology, computational psychology, etc. Such phenomena occur below the threshold of conscious awareness and deal in concepts that are largely alien to our everyday experience and discourse (fantasies about future societies adopting 'brain state language' notwithstanding (Churchland 1981)), though they may borrow certain concepts from personal level discourse in a more-or-less metaphorical way (Dennett 1971). 'Representation' in this sub-personal level sense pertains to information transformation occurring in particular parts of the brain (or elsewhere, cf. chapter IV) which causally contributes to the production of personal-level phenomena like consciousness and belief.

When Merleau-Ponty and others following him (e.g. Romdenh-Romluc 2011a, Dreyfus 2000; 1992) express hostility to the 'intellectualist' thesis that in action, one's bodily movements must be the object(s) of a mental representation, it is primarily the personal level sense of 'representation' they have in mind. In other words, in following Merleau-Ponty in denying that embodied subjects must represent their body to themselves in order to engage in absorbed coping, I am denying that the attribution of a object-intentional mental state about their body is required to explain their behaviour and/or to describe their lived experience. As both a phenomenologist and a theorist writing before the discipline of psychology had entered its current cognitivist paradigm, Merleau-Ponty simply has no interest in the sub-personal version of representationalism (though one might reasonably wonder, as does

Dreyfus (1992), whether the inappropriateness of representational talk at the personal level has ramifications for a theory of cognition which posits corresponding representations at the sub-personal level of cognitive architecture). Similarly, Gallagher's claim that the body schema is not to be identified with, or reduced to, a mental representation of the body, is often advanced as a phenomenologically-motivated personal-level claim. Neither philosophers' views are strictly incompatible with the idea that having a functioning body schema requires the presence of certain kinds of sub-personal neural representations. They are, however, incompatible with the claims that such representations constitute a continually present body image, and that the body schema is identical or reducible to these neural representations (Gallagher 2005 p.134).

Though the ability to deploy the appropriate motor programs required to perform a given action or sequence of actions will involve the activation of certain neural representations, this does not change the fact the acting subject need not, generally does not, and in many cases cannot, represent these motor commands to themselves in thought. At the personal level of description, their experience is harmonious with the project at hand and the content of their mental state far less fine-grained than that of their sub-personal states. Generally, both when making first-hand reports of what we are doing and describing the observed actions of others, we describe our intentions at the 'highest pragmatic level of description', making recourse to norms of behaviour rather than rational reconstructions thereof (Gallagher 2005 p.238, Ratcliffe 2007). If I am interrupted while reaching for my coffee cup and asked what I'm doing I might respond 'picking up my cup to drink from it', just 'picking up my cup' or, more likely, simply 'having a drink', but not 'I am computing the trajectory of my fingertips, hand and arm using visual information and prorioceptive feedback in

order that motor signals from my premotor cortex might initiate the following

sequence of finger movements and postural adjustments via electrochemical signals

sent to my central nervous system ...' (or whatever). The implicated neural

representations, then, though perfectly physically real, are not strictly *mental*

representations - they do not feature in the agent's psychology - though they may

form part of the physical substrate of a mental representation which does - a body

image. For instance, it is often assumed that some form of cognitive access to a

neural body map is required for subjects to be able to conceive of their bodies as

their own and to possess a concept of themselves as an embodied subject (e.g.

Metzinger 2009; Damasio 2000). Ramachandran, for example, hypothesises that

neuropathological conditions like anosognosia (inability to acknowledge a disability,

such as paralysis of a limb) and somatoparaphrenia (delusional denial of ownership

of a body part) are explicable in terms of damage to the cortical pathways required

for such access (Ramachandran & Blakeslee 2005 p.253-4). But qua 'self *concept*'

(ibid., my italics), the body map here serves as the physical substrate, or part of the

physical substrate, of the subject's body *image*. In our ordinary pre-reflective being-

in-the-world this conceptualised body image, though cognitively accessible upon

reflection, is recessive to the point of absence. We can engage in absorbed coping

without thinking about our bodies at all, because the need for such a mental state is

alleviated by the successful close-to-automatic functioning of the body schema itself.

So while it must be admitted that the body schema and the body image are likely to

have some overlap at the sub-personal level of cognitive architecture, they are

nonetheless conceptually distinct for this.

As a system of bodily dispositions dependent upon sub-personal motor processes, the

body schema straddles the personal/sub-personal distinction. What then, *is* the

relationship between sub-personal neural representations and the body schema? The answer is simple: certain neural representations are *part of* the physical processes which comprise the body schema. Here Clark's (1997) notion of an 'action-oriented representation' - which I have already smuggled into the chapter via the quote from de Vignemont - is instructive. Talk of action-oriented representations is increasingly common in neuroscientific treatments of the inner body map, though Clark coins the term in the context of ecological robotics; robots capable of coordinated behaviours in the absence of pre-programmed rules or inner models of their environments (see e.g. Brooks 1991a; 1991b). These artificial creatures are capable of negotiating their way around a laboratory whilst avoiding obstacles (including people trying to distract them), and collecting discarded coke cans (ibid.). Action-oriented representations are sub-personal information-carrying states which 'simultaneously describe aspects of the world and prescribe possible actions, and are poised between pure control structures and passive representations of external reality' (Clark 1997 p.49). As Wheeler insightfully points out, the content of these minimal representations and their behavioural effects are dependent upon the context in which they are active; what they mean depends on what the system in which they're embedded is doing (Wheeler 2005 p.196). In other words, their intentionality is operative rather than objective. Sub-personal representations of this kind are therefore good candidates for the sorts of neural representations implicated in pre-reflective absorbed coping. Acting in the world requires modulation on both sides of the boundary of the body. As I shall argue in the next chapter, we respond skilfully to the ways in which the world presents itself to us in perceptual experience, and the morphology of the body provides the skill set by which we respond to it accordingly. Just as a body sans action-oriented neural representations cannot act upon the world, action-oriented

representations mean nothing, and do nothing, outside of the practical context in which they are embedded.. The content of action-oriented representations is therefore 'wide' (determined by their environment) rather than 'narrow' (determined by internal properties of the system) (see McCulloch 2003, Putnam 1975, Burge 1979), and relatively transient compared to say, a belief. As Wheeler also notes, context-sensitive action-oriented representations are precisely the sorts of inner features one would expect to find in organisms who have evolved to enjoy the kind of primordial being-in-the-world envisioned by Heidegger (and by extension, Merleau-Ponty) (ibid. p.198-9).

6. Plastic brains and rubber hands

Thus far I have implicitly treated the body schema as an integrated collection of functions of the human organism. In this section I present an array of cases whereby paraphernalia external to the boundaries of the biological body may be considered to be incorporated into the body schema. This theme of bodily extension will be revisited in chapter III.

There is ample empirical evidence to support the claim that the neural representations underlying the functions of the body schema exhibit a considerable degree of plasticity and context sensitivity, thus justifying the claim that they are action-oriented. Head and Holmes (1911) were taken by the way in which a blind man using a stick to negotiate his surroundings localised tactile sensations at the tip of the cane rather than in his hand, and hypothesised accordingly that his brain must treat the cane as part of his arm (Merleau-Ponty also notes this phenomenon (Merleau-Ponty 2012 pp.144-5)). It turns out that they were correct; the brain's

representation(s) of the limbs is highly adaptable and continually modified by visual, tactile and kinaesthetic feedback. Macaque monkeys can be trained to use a rake to retrieve food pellets. During use of the rake, the receptive field[9] of neurons in their somatosensory cortex - the cortical location of the inner body map -extended beyond the boundary of the monkey's arm to incorporate the area occupied by the rake itself (Iriki, Tanaka & Iwamura 1996). Similar effects have been observed in human tool use also (Maravita & Iriki 2004, Maravita, Spence & Driver 2003). While using a grasping tool (a bit like the device used by park keepers to pick up litter) to estimate the size of an object, the brain treats the source of the haptic information (the signal from the hand) as originating from the tip of the tool rather than the actual hand (Takahashi, Diedrichsen & Watt 2009). Following this kind of tool use, subjects are temporarily prone to overestimate the distance between tactile sensations on their arm, suggesting that, for a short time, their brains continue to represent their arm length as still extended (Cardinali et al 2009).

The so-called 'rubber hand illusion' provides another example of incorporation of extra-biological items into the neural body map. Botvinick and Cohen's (1998) subjects' sat at a table with their left hand hidden from view with a life-sized rubber hand placed in front of them in an anatomically plausible way. Two small paintbrushes were used to stoke both the occluded real hand and the perceived rubber hand simultaneously. After sitting with their eyes fixed on the fake hand for ten minutes, many of Botvinick and Cohen's subjects reported feeling what Tsakaris, Prabhu & Haggard (2006) have since dubbed 'proprioceptive drift' - a sense of

---

[9] The area to which an applied stimulus - usually a bodily surface - will activate the firing of particular neurons.

ownership for the rubber hand, and claim to feel tactile sensations in it, rather than their hidden real hand. The effect was tested against a control group where the stroking of the rubber hand was asynchronous with the stoking of the real hand. The number of subjects who claimed to experience the illusion dropped to seven per cent, as opposed to forty-two percent in the synchronous condition, indicating that synchronisation between tactile stimulation and visual perception of the perceived stimulus is important to the creation of the illusion (ibid. p.756).

Armel and Ramachandran (2003) replicated Botvinick and Cohen's experimental setup and measured the galvanic skin conductance response[10] of the subjects who experienced the illusion. Once the illusion set in, Armel and Ramachandran 'injured' the rubber hand by bending back one of its fingers. This elicited a strong skin conductance response in their subjects and they exhibited pain-anticipating avoidance behaviours like wincing and even snatching their real hand away from experimenters (ibid. pp.1502-3). It should be noted, however, that only two out of one-hundred and twenty subjects reported feeling any actual pain (ibid. p.1503). Subsequent studies reveal that the rubber hand illusion admits of degrees of strength and its onset proceeds through several stages before felt ownership of the rubber hand can displace that of the real hand (Moguillansky, O'Regan & Petitmengin 2013). While the skin colour of the rubber hand apparently has no effect on subjects' susceptibility to the illusion (Farmer, Tajadura-Jiménez & Tsakiris 2012), the strength of the illusion is inversely proportional to the distance of the rubber hand from the real hand (Loyd 2007), and depends on the prop's being suitably hand-like in appearance; synchronous stroking of the hidden hand with a plank of wood will

---

[10] An electrochemical response to stimuli, namely the release of sweat. Measuring galvanic skin conductance is a popular method of testing the strength of emotional responses to stimuli.

not give rise to the illusion, a coarsely hand-shaped piece of wood might, and an anatomically accurate model hand probably will (Tsakiris et al 2010).

Cases like these straddle the schema/image distinction. Neural adaptation to tools occurs in action contexts approaching absorbed coping, which under Gallagher's criteria would constitute a modification to the body schema, but rely on conscious visual monitoring of the artificial body part, which involve aspects classifiable under body image. Similarly the rubber hand illusion consists in the temporary modification of proprioception (schema) as the product of visual attention to a fake limb (image). While it is highly unlikely that subjects in the grip of the rubber hand illusion would conceptualise their bodily selves as incorporating the fake hand, it seems that a certain degree of conceptual knowledge constrains the illusion, as anatomically implausible props won't induce an illusion of ownership. At the same time, the relationship between spatial proximity of the fake hand to the real hand suggests that proprioceptive norms constrain the illusion also. We must therefore concede that the schema/image distinction is heuristic rather than exhaustive, and acknowledge interaction between the two. I do not think that the distinction is in too much trouble, though. It still provides a useful conceptual tool for distinguishing aspects of our mental lives - object-intentional mental states about our bodies, and operative-intentional absorbed coping. And, as the above remarks (admittedly, sketchily) indicate, it can be fruitfully deployed to analyse the different stages of the illusion and their phenomenal and functional characteristics.

In data-driven discussions of these neural adaptations, it is generally assumed that the body schema is identical or reducible to the continually-updated neural body map. Metzinger, for example, invokes both Iriki et al's macaque experiment and the rubber hand illusion to support the thesis that embodiment is entirely a matter of what's

going on in the brain (Metzinger 2009 pp.75-82). I think this is too hasty. For one thing, we should be wary of drawing any conclusions based on data from skin conductance alone because, as de Vignemont (2011) points out, galvanic skin response does not suffice to establish bodily incorporation. The same sort of response occurs when we see other people in pain (Westbury & Neummann 2008) and even hardened viewers might sweat, wince or grimace when watching an especially gruesome scene from a horror film, so this kind of response probably has more to do with empathy than a modified body schema as such (see e.g. G.Hein et al 2011). This is consistent with the fact that people with autistic spectrum disorders are less susceptible to the illusion (Palmer at al 2013, Casico et al 2012).[11] Recall also that Armel & Ramachandran's recoiling subjects snatched away their real arm, thereby exercising a retained capacity for a body-schematic response to a perceived situation. Habits die harder than transitory illusions.

As for proprioceptive drift itself, visual processing has a tendency to dominate or overrule that of other senses. In the McGurk effect, for example visual perception of a speaker's face overrides audition, misleading people into hearing syllables different to those being said (McGurk & MacDonald 1976)[12]. Given this fact, it is perhaps not too surprising that this should be the case with proprioception also (Stokes & Biggs forthcoming). But it doesn't show that an internally-constructed body map magically erases the body schema proper. The rubber hand illusion is relatively short-lived and

---

[11]To the best of my knowledge, there are currently no data on skin conductance from autistic subjects subjected to the rubber hand illusion. Going out on a limb (so to speak), I would venture the prediction that they would show a reduced skin conductance response.

[12] When shown a video of a face uttering the syllable 'ba', but with some of the 'bas' overdubbed with the syllable 'ga', people hear the syllable as 'da' (McGurk & MacDonald 1976).

is induced in artificial experimental conditions where the role of the body is conceived as entirely passive. Experiments with video projections of fake hands strongly suggest that achieving a complete sense of agency, as opposed to mere temporary ownership of the hand[13], and to fully integrate the prop hand into the subject's proprioceptive field, requires active manipulation of the prop (Tsakiris, Schutz-Bosbach & Gallagher 2007, Tsakaris, Prabhu & Haggard 2006), because the sense of bodily agency requires interplay between centralised brain processes and afferent input from the peripheral nervous system (Tsakaris & Haggard 2005; Marcel 2003). The rubber hand illusion therefore presents a case of partial integration, or an early stage of an as-yet incomplete integration, of a foreign object into the body schema, A more thorough integration requires a temporally extended process of active bodily engagement.

7. Embodiment ain't in the head

Neural plasticity allows for flexible embodiment, but it is only part of the story. Some philosophers have defended the view that proprioception provides a kind of direct knowledge of one's body that is immune to error (e.g. O'Shaughnessy 1995, Evans 1982). Prima facie, the rubber hand illusion and data from tool use experiments contradict this thesis (cf. de Vignemont 2014). Considered phenomenologically, however, this is not quite true. The greater degree to which an object integrates with the body schema via its utilisation in action, the more inappropriate the term 'illusion' becomes to characterise the experience. It is of

---

[13] Bodily ownership is the sense of 'mineness' about the body, that *this* body is *my* body. The sense of bodily agency is the awareness that an action is something one is *doing* (see Gallagher 2000).

course true that if, say, Head and Holmes' blind man were to affirm the belief that his hand is located several feet away from the end of his wrist he would be saying something false. But viewed phenomenologically, the experience of displaced touch and proprioceptive drift constitutes an extension of the body through which the blind man gains unmediated access to features of his surroundings (e.g. that a rock is in front of his feet). Here the aforementioned distinction between the objective and the lived/phenomenal body is pertinent. Insofar as the object is utilised in an active motor project, it extends the lived body beyond the boundaries of the objective, biological body. As Merleau-Ponty remarks on this example

> The blind man's cane has ceased to be an object for him, it is no longer perceived for itself; rather, the cane's furthest point is transformed into a sensitive zone, it increases the scope and the radius of the act of touching and has become analogous to a gaze. (Merleau-Ponty 2012 p.144)

Similarly, amputees using prosthetic hands can hone new motor skills through use, sometimes performing as well on tasks (e.g. accurate grasping) as their able bodied counterparts (Schabowsky et al 2008) and experience the end point of a surviving partial limb as coinciding more with the prosthesis than their stump (McDonnell et al 1989).

Doing justice to incorporation requires taking the phenomenology seriously (as, for example, does Clark 2003). In the spirit of naturalism and physicalism, we should readily accept that the phenomenology of embodiment is ultimately constituted by physical states and processes. But while it is beyond reasonable doubt that there exist sub-personal neural representations of body parts and/or non body parts which are poised to govern action, it does not follow, in a reductionist spirit, that 'the body

49

schema *is* [...] as a cluster of sensorimotor representations that are action-oriented'
(de Vignemont 2010 p.679, my italics). As I seek to articulate in more depth over the
next two chapters, our perceptual encounters with the world involve an implicit
understanding of, and ability to respond to, items, events and situations which are
grasped as *meaningful* for us. We see objects as 'affording' possibilities for action
(Gibson 1986), and our perceived surroundings solicit what Merleau-Ponty calls
'motor significations' which have meaning neither for a detached Cartesian ego or a
group of firing neurons but, as Merleau-Ponty, following Heidegger, stresses, a
bodily subject inescapably embedded in a world of practical engagements. Intelligent
behaviour therefore comprises a 'sensory circuit' (Merleau-Ponty 2012 p.89) between
the active body and the perceived world, of which action-oriented representations are
merely one small but vital physical constituent. So while a body schema cannot
function in the absence of certain sub-personal action-oriented neural
representations, the concept of a body schema does not signify these neural
representations, or any singular neural representation they might be said to jointly
constitute, but the network of bodily dispositions itself, which owe not just to the
inner workings of parts of a subject's brain but also their bodily morphology and skill
set.

8. A Note on Phantom Limbs

Phantom limbs present a limit case for a phenomenology of embodiment. Amputees
often experience painful, tactile or kinaesthetic sensation emanating from their absent
body parts. There are also documented cases of aplastic phantoms, phantom limb
experiences in those born without the corresponding body part (usually a limb) (Price

2006), which are likely to emerge from a combination of genetic and developmental factors (ibid.). It is tricky to say anything definite about phantom limbs as their phenomenal characteristics, physical causes, and psychological histories vary so much between cases (for a detailed empirical review, see Giummarra et al 2007). Nevertheless, I am obliged to comment on them, albeit cautiously, because they present a prima facie challenge to my claim that embodiment is not reducible to events in the brain.

Though some patients experience their phantoms as immobile, others report being able to move them intentionally, and these movements can constrain the movement of surviving body parts in the same way as real limbs despite the absence of visual or proprioceptive feedback (Franz & Ramachandran 1998). In the neuroscientific and clinical literature, phantom limbs are often described in terms of a body image, but authors typically have in mind aspects of a body schema also (Gallagher 2005 ch. IV). Insofar as phantoms show up in patients' consciousness - as phenomenal curiosities, obstacles to an ordinary life, or sources of pain or discomfort, they are best characterised in terms of body image - an occurrent intentional object of a mental state. But insofar as they depend on the survival of the neural representations underlying proprioception and kinaesthesia, they are the product of a partial body schema (ibid.).

Phantom limbs are adaptable, and the ways in which they are experienced are highly amenable to perceptual situations and contexts of action. Phantom movement is experienced in some contexts but not others. For example, one patient who was born without arms feels two short phantoms which gesticulate in conversation, but seem to slump motionlessly while walking (Ramachandran & Blakeslee 2005 p.40-2). Immobile phantoms (i.e. where tactile or pain sensation are felt but no movement)

can be revived using mirrors to invoke the visual appearance of the subject's intact hand taking the place of their missing hand (Ramachandran, Rogers-Ramachandran & Cobb 1995). This works by exploiting visual dominance and fooling the patient's brain into attributing the outcome of an intention to act to the perceived movement, and can be used to ease spasms felt in the missing hand (ibid.). The newly acquired felt movement is transient, however, and is lost without visual perception of the mirrored superimposed limb (Ramachandran & Blakeslee 2005 p.53). Conversely, phantom movement can be eliminated through 'learned paralysis', where the lack of visually perceived movement nullifies the subject's act intentions (ibid. p.46). Amputees who use prostheses tend not to feel phantom pains, but come to feel their phantom as being embodied by the prosthetic limb (Giummarra et al 2007 pp.222-3), and it is thought by some doctors that a phantom is essential to successful use of a prosthetic limb (Sacks 1985 p.71). Following adaptation, they can sometimes even scratch phantom itches by scratching their prosthetic limb (Giummarra et al 2008 p.154).

Despite their being experienced as private mental episodes, the surviving neural representations of absent body parts apparently depend heavily for their content on perception and action. This is consistent with my characterisation of the neural representations underlying the body schema as action-oriented. Though the experience of phantoms owes initially to preserved sub-personal neural representations of missing body parts, the content of these representations is wide rather than narrow; what they represent, and what shows up in the subject's experience on their basis, is determined by the perception of external events and the subject's intentional action. The adaptability of phantom limbs suggest that they depend in large part upon factors external to the brain, and that the subject's

(partially) illusory embodiment is constituted, at least in part, by their active

perceptual engagement with their environment. Even in the absence of a real limb,

the internal body map is malleable and depends for is content on the ways in which it

is actively plugged into the world.

II. Embodying Perception: The Sensorimotor Approach to Vision

1. Introduction

In the previous chapter, I introduced and defended a distinction between body schema and body image, and argued that the former is not reducible or identical to a sub-personal neural representation. In this chapter I explore the ways in which embodiment structures perception, with a particular focus on visual experience, by means of a discussion and defence of the sensorimotor approach to visual perception as formulated in the empirical work of Kevin O'Regan, the more theoretical-philosophical work of Alva Noë, and their collaborative efforts. In brief, the sensorimotor approach characterises visual perception as the active exploration of one's immediate environment via the exercise of a special kind of bodily skill.

I begin with a discussion of the position the sensorimotor approach attacks, the so-called 'snapshot' conception of vision, which comes in both personal and sub-personal variations. The case against the snapshot conception is illustrated by means of experimental work on change blindness and inattentional blindness, and phenomenological analysis courtesy of Noë, Husserl and Merleau-Ponty. I then turn to the concept of 'sensorimotor understanding', which is of central importance to the sensorimotor approach. I argue that Noë's account suffers from certain contradictions resulting from some overly strong but ultimately non-essential claims about the role of action in visual experience. These shortcomings are rectified by situating Noë's account in relation to Husserlian phenomenology, then appealing to particular aspects of Merleau-Ponty's philosophy to amend Noë's account. The chapter concludes with a response to Andy Clark's critique of the sensorimotor approach.

2. Two versions of the sensorimotor approach

Although the sensorimotor approach was brought into the philosophical and scientific mainstream by O'Regan and Noë's (2001) seminal *Behavioural and Brain Sciences* target article, the two former collaborators have since appropriated the sensorimotor approach to quite different philosophical ends. As Bishop and Martin (2014) instructively highlight, O'Regan takes the sensorimotor approach, supplemented by a higher-order thought theory of self-consciousness (a la Rosenthal 1997), to provide the necessary conceptual tools to solve what Chalmers (1996; 1995) famously dubbed the 'hard problem' as opposed to the 'easy problems' of consciousness (namely, the perplexing metaphysical conundrum of why any physical system should give rise to conscious experience *at all*) in an empirically testable and naturalistic way. This is how O'Regan presents his project throughout his latest book, *Why Red Doesn't Sound Like a Bell* (O'Regan 2011) and elsewhere (O'Regan 2010), the consequences of which he takes to be instructive with respect to the possibility of machine consciousness (O'Regan 2012). Noë, on the other hand, bar a few sketchy asides, is less worried about the hard problem of consciousness or the possibility of conscious robots and repeatedly asserts that his philosophical project is essentially a phenomenological one (e.g. Noë 2004 p.33; 176). That is to say, he is more concerned with the ways in which action and embodiment structure the content of perceptual experience. *This* project, Noë readily admits, necessitates 'taking a little bit of consciousness for granted' (ibid. p.230, italics removed). Though I am generally sympathetic to O'Regan's take on the issue of the hard problem (though less so on machine consciousness), and will of necessity discuss a good deal of his empirical and theoretical work as well as the aforementioned collaborative paper, I

am primarily concerned with the sensorimotor approach as it functions in Noë's project. That it is, I am concerned with building on my previous chapter in order to explore the way in which embodiment shapes our perceptual experience.

A quick note on terminology. The sensorimotor approach is often also referred to as the 'enactive' approach to vision. At one point, this was Noë's preferred label (i.e. in Noë (2004), cf. Noë (2012)). The term 'enactive' was introduced into contemporary philosophical and cognitive-scientific vocabulary by Varela, Thompson and Rosch (1991) in *The Embodied Mind,* which synthesises a novel blend of phenomenology and Buddhist thought to form the basis for a critique of what they take to be a latent and problematic Cartesianism inherent in cognitive science (see also Thompson 2007). In recent years, however, the label 'enactivism' has come to be synonymous with the rejection of any theory of cognition which invokes representational or informational content (e.g. Hutto and Myin 2013, Hutto 2013). In order to avoid conflation with these and other views presented under the banner of enactivism, and for further technical (though slightly pedantic) reasons explained later in this chapter, I shall adhere to O'Regan and Noë's (2001) original terminology throughout.

3. Against pictures in the head: the 'snapshot' conception of vision

In his neo-classic book *Consciousness Explained,* Daniel Dennett (1991) rages against a position he dubs 'Cartesian materialism', which he take to have had a deeply unfortunate and regrettable influence on a good deal of philosophical and scientific theorising about the nature of the mind. Roughly, Cartesian materialism is the belief in a particular centre of the brain where consciousness *happens*, where what we decide, perceive, remember, feel or imagine is presented to a centralised

subject housed inside an interior mental realm - the 'Cartesian theatre'. Dennett argues, very convincingly, that the Cartesian theatre is a myth and, as an instance of the homunculus fallacy, explains nothing. The sensorimotor approach to perception rebels against one particular instance of Cartesian materialism, namely what Noë helpfully dubs the 'snapshot' conception of vision (Noë 2004 p.35). According to the snapshot conception, visual perception is essentially pictorial in nature. Although Noë himself does not always distinguish clearly between them, there are in fact two related versions of the snapshot conception of vision. One version applies at the personal level of visual phenomenology, while the other applies at the sub-personal level of the cognitive processing which underpins that phenomenology. The personal level version of the snapshot conception has it that visual experience is comparable to a richly detailed picture, like a photograph, of the external world. On this conception of perceptual experience, vision presents us with rich images of the world the vivid detail of which is given instantaneously and simultaneously to the perceiver. The sub-personal version of the snapshot conception has it that visual phenomenology is caused by, supervenes on, or can be identified with, states of the visual system which encode stable, detailed representations of the external world constructed out of sensory inputs. I shall consider the sub-personal version of the snapshot conception first.

Aside from certain metaphors inherited from early modern philosophy,[14] much of the motivation for the snapshot conception of perceptual cognition stems from considerations about the poverty of the eye. Retinal images, as images go, are pretty poor. As O'Regan is fond of pointing out, the lens of the human eye is vastly inferior to those found in even moderately priced cameras. Any photographs made with a

---

[14] Namely the 'camera obscura' metaphor for perception generally attributed to Locke, see e.g. Bailey (1989).

lens with as low a resolution as the one found in the eye would be blurry almost all over bar the very centre, and even there the colours would be vague and muddled (O'Regan 2011 p.10). The photoreceptors that allow us to discriminate colours are found only in the centre of the eye, rendering all but the centre of our visual field virtually colour blind. Each eye has a vascular scotoma, a blind spot caused by blood vessels between the optic nerve and the retina which obstruct the images formed on the retina. On top of this, the brain has two different retinal images to deal with. These are only two-dimensional, and they are upside down (for a short but thorough account of the shortcomings of the human eye, see O'Regan (2011) pp.3-22). Furthermore, the eyes saccade on average around three times per second (Yarbus 1967) and we blink several times a minute, which renders individual retinal impressions very short lived. In spite of all these optical defects, we perceive a stable and detailed three dimensional world incredibly well. Mainstream cognitive science has therefore by and large conceived of the scientific study of vision as the task of accounting for the mechanisms by which the brain forms internal models of the world out of paltry sensory stimulation. For example, Marr's landmark book *Vision* criticises Gibson's (1986) competing 'ecological' theory of vision as the 'direct pickup' of information from reflected light for allegedly underestimated the complexity of the information processing required by visual cognition (Marr 1982 p.253), and lays the foundations for contemporary computational vision science by treating visual perception as a multi-stage process by which initial images are processed to form a three dimensional model of what is perceived (see ibid. p.260 for a schematic diagram of the stages of processing).

Sensorimotor theorists like O'Regan and Noë question just how much detail need be encoded in these internal representations. Their answer, in contrast to their

predecessors in vision science, is 'not very much'. Evidence for this hypothesis comes by way of various experiments on the related phenomena of change blindness and inattentional blindness, the results of which were originally predicted by Dennett (Dennett 1991 pp.467-8). Change blindness occurs when subjects fail to notice changes in a visual stimulus despite looking straight at it. Grimes (1996) conducted a now famous experiment were subjects fitted with eye tracking devices were instructed to read a page of text displayed on a screen. Unbeknownst to the test subjects, text surrounding the small number of words upon which they could focus at any one moment was constantly changing, but they did not notice. Only the handful of words immediately present to the fovea can be read in one go, rendering changes to text in the periphery of the visual field undetectable (ibid.). O'Regan and colleagues (2000) used eye tracking technology to monitor subjects' blinks while they were presented with a series of pictures. While they blinked, changes were made to the picture - objects were added or removed, colours were changed, shadows altered, etc. Again, subjects consistently failed to detect the changes. I have personally seen O'Regan demonstrate this effect outside of the lab and without the aid of eye tracking, simply by showing a conference audience a short video, cleverly disguised as a photograph, of a Paris street scene. The colour of a car displayed prominently near the middle of the 'photograph' was gradually altered from blue to red, but the vast majority of the audience failed to notice this (and those who did notice were all familiar with O'Regan's work).[15]

Similarly, inattentional blindness occurs when subjects attending to part of a visual scene fail to notice an otherwise striking event elsewhere in that scene. Simons and

---

[15] A video of Noë replicating O'Regan's trick can be found here (at around the six and a half minute mark): https://www.youtube.com/watch?v=af3Vq-C1ck8 (last accessed: 8/8/14).

Chabris' (1999) well known 'gorilla suit' experiment provides perhaps the classic example of inattentional blindness. Subjects are shown a video of a small group of people playing catch with two basketballs. Half of the people are wearing blue jeans and white t-shirts and the other half are clad entirely in black. Some test subjects are instructed to count the number of passes made by the black team and the others are instructed to count the passes made by the white team (each team has their own ball). During the course of play, a man wearing a gorilla suit strolls into the midst of the players, beats his chest, then leaves. The subjects asked to count the black team's passes immediately notice the gorilla. But up to sixty percent of the subjects tasked with counting the white teams' passes do not notice him at all (ibid.). Those tracking the black teams' passes are looking out for black-clad players, so the black gorilla suit grabs their attention. Those tracking the white teams' passes are attending to white-clad players at the expense of darker parts of the visual scene, so are prone to miss the giant ape man.[16] O'Regan also cites Haines' (1991) remarkable example of an experienced airline pilot who, in what was thankfully only a NASA training simulation, landed on top of another plane parked in full view on the runway while attending to the visual display on his windshield (O'Regan 2011 pp.50). As O'Regan also notes (ibid. p.51), many ordinary traffic accidents are the direct result of drivers 'looking but failing to see' cyclists, pedestrians and even parked cars (for a review see Langham et al 2002).

Although it is of course true that our attention is often draw to fluctuations in the periphery of the visual field (think of the experience of suddenly noticing a spider running across the floor in the corner of a room), we do not see the exact cause of these events until we turn our head and/or eyes to inspect it; we see that something is

---

[16] Simons and Chabris' original video can be found here:
https://www.youtube.com/watch?v=vJG698U2Mvo (last accessed 8/8/14).

happening without yet seeing what that something actually is (O'Regan tends to refer to this as the 'grabbiness' of vision). Hence, the moral of change blindness and inattentional blindness, according to sensorimotor theorists (and indeed to some non sensorimotor theorists, such as Dennett), is that perceivers do *not* internally represent the features of a visual scene in very much detail at all (I address objections to this interpretation below). The proven poverty of parafoveal vision suggests that at any given instant, the visual system only samples a very small amount of the available information in a visual scene. Perceptual constancy - the enduring stability of the perceived world - therefore owes not to a detailed mental model of the environment, but to the stable layout of the external world itself; the world looks stable to us simply because it is stable! For this reason, O'Regan characterises the role of the external world in visual cognition not as a source of stimulation from which an internal model is constructed which maps the world outside the head, but as a source of 'outside memory' readily accessible to perceivers via selective visual attention (O'Regan 1992). The experience of the stable and continual presence of the visual world owes, on the sensorimotor account, not to cognitive access to a detailed internal model of the world, but to perceivers' 'practical knowledge about currently possible behaviours and associated sensory consequences' (O'Regan & Noë 2001 p.946), that is, in their knowing how to access information in the external world itself.

It is important to note, however, that unlike ecological vision theorists inspired by Gibson (1986) (e.g. Chemero 2009) as well as certain others (e.g. Hutto & Myin 2013), sensorimotor theorists are not completely opposed to visual representations in the more general sense in which psychologists typically use the term, that is, sub-personal states the function of which is to detect and track certain environmental

features, but they are opposed to the widespread assumption that these representations serve as detailed three-dimensional models of the external world. They also deny representations in the more metaphysical sense of a mediating mental entity between perceiving subjects and the distal causes of their perceptions[17]

> What the [sensorimotor theory] denies is not the general claim that there are representations operative in one or various stages of visual processing but that the feeling of seeing all the detail at any moment is the result of a fully detailed, continuously present representation of all the detail. (O'Regan & Myin 2009 p.189)

A line of objection to the invocation of change/inattentional blindness against the sub-personal version of the snapshot conception is pursued by Prinz (2006) and Jacob (2006), both of whom argue that change/inattentional blindness does not suffice to refute pictorialism. Prinz cites experiments performed by Mitroff, Simons and Levin (2004) who found that subjects can fail to consciously notice the differences in two sequentially presented arrays of objects but still perform above chance when forced to indicate, via answers to closed questions, which object was absent from the second stimulus. Hence, Prinz writes

> Another possibility [...] is that the visual system generates very rich representations, but doesn't compare every detail. On the [sic] view, the visual system registers changes in the scene, but not the fact that the change has taken place. The differences are all encoded, but no comparison system picks up on them. This interpretation is consistent with the pictorial theory of

---

[17] Noë tends to flip-flop on the role of representation in vision, sometimes appearing to endorse anti-representationalism across the board but at other times urging for a more liberal notion of representation without abandoning it outright, sometimes in the space of the same paper. See e.g. Noë (2010).

vision, and Noë [and O'Regan] must argue against it to show that [their] interpretation is right. (Prinz 2006 p.13)

Here Prinz begs the question against the sensorimotor view's anti-pictorialist stance by first assuming the existence of detailed internal representations of the two visual scenes, then using these to ground his story about a lack of comparison mechanism to detect their differences in content. In order to avoid begging the question, Prinz's argument has to assume that Mitroff, Simons and Levin's findings demonstrate the existence of such representations. But they do not. Mitroff, Simons and Levin are totally upfront about the fact that 'the results do not speak directly to the *nature* or *completeness* of the internal representations' (Mitroff, Simons and Levin 2004 p.1279, italics in original) and that '[t]he present experiments were not designed to differentiate these two alternatives' (ibid.). They even conclude that 'it is important to note that a comparison failure explanation does not rely upon *complete* representations [...] the representations need only to be sufficiently detailed to support change detection' (ibid., my italics), which undermines Prinz's use of them as a premise for his far stronger pictorial version of their 'absence of comparison' account. Furthermore, Mitroff, Simons and Levin's findings do not directly contradict the absence of awareness and subsequent inability to recall visual detail reported by the subjects of Grimes', Simons and Chabris', O'Regan et al's, and Haines' experiments, though admittedly, the extent of their recall upon sufficient prompting is an open empirical question. But it is far too big a jump from the fact that a small amount of detail can be preserved in memory from unconscious visual processing to the claim that the brain encodes detailed 'photographs' of the external world.

Similarly, Jacob, who characterises vision in quasi-snapshot terms as '*recording* a fact' (Jacob 2006 p.2, my italics), reiterating speculation of Simons and Rensinik's (2005), argues that the result of change/inattentional blindness experiments are not strictly incompatible with the existence of detailed internal mental representations

> (a) The detailed representation of the display perceived at t1 could decay or fade away before the comparison takes place at t2. (b) The detailed representation [...] could be encoded in a neural pathway unavailable for comparison. (c) The content of the detailed representation of the stimulus perceived at t1 could be encoded in a format unsuitable for comparison. (d.) Although the detailed representation has not decayed until the comparison takes place and although the representation is in the right pathway and in the right format, still the comparison process could fail for some other reason. (Jacob 2006 p.4)

The sensorimotor theorist can agree that change/inattentional blindness is not, technically, strictly incompatible with rich pictorialism, but given that there demonstrably *is* plenty of detail available from a visual scene at any one time that is neither consciously perceived nor accessible to memory (Grimes' peripheral text, Simons and Chabris' gorilla man, Haines' aeroplane) they would be quite justified in dismissing Jacob's list of detailed representation-saving counterfactuals as ad hoc. Whichever way the data are spun, the visual world does not get to be inside the head.

4. The visual world is not a grand illusion

Some philosophers and cognitive scientists impressed by the results of change/inattentional blindness experiments draw the conclusion that in our pre-

scientific thinking about perception, we are considerably and systematically mistaken about the nature of our own experience. Blakemore and colleagues (1995), following Dennett (1991), suggest that when we take ourselves to enjoy richly detailed visual experience, and thereby a richly detailed visual field, we are subject to an illusion. Despite how things may seem to us, it is claimed, science teaches us that consciousness does not in fact acquaint us with a stable visual world of richly coloured detail, but short-lived and vague visual fragments which are, in Blakemore's charming turn of phrase, continually 'washed away' by eye saccades (Blakemore et al p.1076). This would be disturbing and startling if it were true; the visual world would be a grand illusion![18] Furthermore, if change/inattentional blindness does indeed show that 'our intuitions about our own visual function are far from useful in understanding the construction of our convincing and stable visual world' (Blakemore et al 1995 p.1080) then it's game over for most phenomenologists.

Fortunately, this interpretation of the data on change/inattentional blindness only follows if we assume the personal level version of the snapshot conception of vision. Though Dennett and Blakemore both reject the sub-personal version of the snapshot conception, they remain wedded, at least implicitly, to the personal level version, which they in turn attribute to ordinary folk. Even O'Regan, whose entire research project is levelled against the sub-personal snapshot conception of perceptual cognition, sometimes falls into the trap of sharing with rival vision theorists the assumption that the personal level version of the snapshot conception constitutes his explanandum as the following quotation, which comes from a chapter entitled 'The Illusion of Seeing Everything', illustrates:

---

[18] See the essays collected in Noë (2002) for exploration of this claim.

We think we perceive things as highly detailed as highly detailed and perfect *all over* the visual field. How are we so sure that we really do perceive those details all over the scene. The answer must be because we know we can answer any question about these details if we ask ourselves about them. But to do this we needn't actually have any information stored in the brain about these details. All we require is to be able to immediately access the information in the world [...] We can then also have the *impression* of perceiving the detail all over. (O'Regan 2011 p.58, first italics in original, second added)

The snapshot conception, as its name implies, conceives of visual experience in terms of hypostasised conscious visual states - static images all the detail of which is given to the subject in one fell swoop. Naturally, people who miss the gorilla man in Simons and Chabris' video or the colour switch in O'Regan's Paris street scene are surprised when their attention is later drawn to the changes. But is this really because they previously took themselves to be enjoying a perception of the basketball players or the street akin to a photographic image? It is doubtful that most people outside of a philosophy classroom would characterise their experience quite like this upon even minimal reflection, because *experience is not like this.* Nobody with working eyes is likely to be alarmed by the fact that peripheral vision is pretty useless for the task of discriminating text, or that they are prone to miss otherwise obvious events if their attention is directed elsewhere. How many time do we 'lose' things - keys, phones, glasses, etc - which suddenly 'appear' right in front of us, having been there the whole time? As Merleau-Ponty stresses throughout *Phenomenology of Perception*, indeterminacy is a *positive* feature of the visual field; it is constitutive of ordinary perceptual experience and not something to be reduced or explained away, as the

surrounding vagueness functions as the background against which the objects of attention can appear (Merleau-Ponty 2012 p7; pp.31-3). Indeed, without this indeterminacy, we cannot even make phenomenological sense of the idea of bringing something to visual attention, 'To pay attention is not merely to clarify some pre-existing givens; rather, it is to realise in them a new articulation by taking them as figures' (ibid. p.32, italics removed). And as he notes elsewhere, it is precisely this phenomenon of the objects of visual attention 'emerging' from a background of indeterminate vagueness that impressionist painters aimed to capture by abandoning the photographic realism of their predecessors (Merleau-Ponty 2004). O'Regan's occasional rhetorical lapses aside, sensorimotor theorists share this Merleau-Pontian sentiment, and emphasise that, contra the snapshot conception, perceptual experience is fundamentally an *exploratory* and *temporally extended* affair; the perceived world is revealed to us not instantaneously, but over time and via a process of perceptual exploration (Silverman 2013). We are constantly moving our eyes, turning our heads, and redirecting our attention, in order to bring things into view. Similarly, when an event in peripheral vision catches our attention, it prompts us to further exercise these abilities. Visual experience is, in this sense, a highly *active* phenomenon, comprising far more than the merely passive reception of sensory input.

The visual field is not restricted by anything analogous to a picture frame. Phenomenologically, there are no clearly discernible borders between where a visual scene begins and ends. Insofar as the visual field is bounded, these boundaries are in continual flux. Though the amount of visible detail discriminable at any one instant may be minimal, our visual experience is always already informed by a sense of the *presence* and *accessibility* of further detail in our immediate surroundings. We do not

typically take ourselves to enjoy all the detail of a visual scene simultaneously, but we do have a implicit sense of what kind of detail is present in our surroundings and available for visual inspection, because as perceivers we know how to access features of a visual scene by tuning our heads, focussing our attention, moving our eyes, interacting with objects, etc. As Noë puts this point, 'The world shows up for us in experience only insofar as we know how to make contact with it, or, to use a different metaphor, only insofar as we are able to bring it into focus' (Noë 2012 p.2). And as Merleau-Ponty observes, the visual field comprises both the indeterminacy of peripheral vision and the knowledge that aspects of objects can be brought into view via movements of the body and of the objects themselves in relation to the perspective imposed by the body

> To say that I have a visual field means that I have an access and an opening to a system of visible beings through my position, and that they are available to my gaze in virtue of a kind of primordial contact and by a gift of nature, without any effort required on my part. (Merleau-Ponty 2012p.224)

As I elaborate below, the 'gift of nature' to which Merleau-Ponty refers is the structure of experience afforded by our embodiment and the way in which the body is situated in its environment. For now though, it suffices to note that the worry that we systematically delude ourselves about the extent of our visual experience is dissolvable once we give up the phenomenologically inept personal level version of the snapshot conception of vision. As Noë puts the point:

> The sceptical reasoning relies on a bad inference from the character of a single visual fixation to the character of seeing itself. From the fact that, when I stare at a point on the wall, I can't see colours in the periphery, it doesn't follow that there are no colours in the periphery of the visual field. For my

visual field, rather, is made available by *looking around.* We look here, then there, and in this way we gain access to the world and our experience acquires that world as content. It is no part of our phenomenological commitments that we take ourselves to have all that detail at hand *in a single fixation.* (Noë 2004 p.57, italics in original)

By paying closer attention to the phenomenology of visual experience we can agree with O'Regan (see quotation above), Dennett (1991 p.354-5), and others, that the brain need not represent all the available detail of a visual scene in its entirety, and that the fact that we experience the visual world as stable and richly detailed is owed to an understanding of the availability of the relevant detail upon inspection, without conceding that we are in any way mistaken about the nature or extent of our own experience, because the content of visual experience is not determined *just* by what impacts upon the retinas, or by what gets represented in the visual cortex, over any single arbitrarily short temporal interval. Clark captures perfectly the way in which the argument from change/inattentional blindness to the 'grand illusion' of the visual world, rests on a sleight of hand, inferring a false phenomenological story from an empirical one from which it does not follow

[T]he Grand Illusion [...] was a trick of the light. [...] The scene before us is indeed rich in colour, depth and detail, just as we take it to be. And we have access to this depth and detail as easily as we have access to facts stored in biological long-term memory. It is just that in the case of the visual scene, retrieval is via visual saccade and exploratory action. Our daily experience only becomes misleading in the context of a host of unwise theoretical moves and commitments: commitments concerning the precise role of internal representations in supporting visual experience, as well as our pervasive

neglect of the cognitive role of temporally extended processes and active exploration. (Clark 2002 pp.201)

Merleau-Ponty gives us a convenient diagnostic term for the fallacy underlying the personal level version of the snapshot conception: the 'experience error'. The experience error consists in the (mis)description of perception in terms of what we come to believe upon reflection about its objects at the expense of capturing the original phenomenon as we live through it (Merleau-Ponty 2012 p.5). It is this fallacy that generates the snapshot conception. Once the fallacy is exposed, we are free to give it up. The visible world transcends any momentary gaze, and seeing is not at all like having a photograph projected inside the Cartesian theatre.

5. The phenomenology of sensorimotor understanding

It is not just objects in peripheral or parafoveal vision that are experienced as present in our ordinary experience. Though strictly speaking un*seen*, *occluded* surfaces and features of objects are likewise given in visual experience as present to one. For example, when I look at a coffee cup on a desk, it is part and parcel of my experience of the cup that it has a reverse side which, though occluded given my current perspective, is present in my immediate environment and potentially visible from an alternative perspective that I could occupy. Phenomenologically speaking, this is just a basic fact about what it is for human beings to see a three-dimensional object as a three-dimensional object; so as long as we are doing phenomenology rather than, say, metaphysics, this feature of experience may be considered irreducible. Husserl revisited these phenomena of perspective and presence throughout the development of his phenomenology. In Husserl's terminology, occluded features of objects are

experienced as 'co-present' with their facing sides (Husserl 1999b p. 222), because the entire perceptual field exhibits the phenomenal feature of *horizons*; any visual presentation of an object at a particular instant anticipates additional presentations of that object at future instants

> [I]n being there itself, the physical thing has for the experiencer an open, indefinite, indeterminately general horizon, comprising what is itself not strictly perceived – a horizon (this is an implicit assumption) that can be opened up by possible experiences. (Husserl 1988 p.23)

We can summarise this Husserlian insight by saying that visual experience has a *horizonal structure.* Contemporary sensorimotor theorists share Husserl's interest in this feature of perceptual experience. Noë's rather confusing term for it in *Action in Perception* is 'virtual presence', but I shall adhere to the original Husserlian terminology. (Nb. To anticipate a potential confusion, it should be noted that 'horizons' enjoys varied usage in the phenomenological tradition. Husserl and Merleau-Ponty sometimes use 'horizons' in an extended sense to encompass not just co-presented features of objects, but also what we would now following Gibson (1986) refer to as 'affordances'. This is unfortunate, but the ambiguity owes more to literary convention than a genuine ignorance of the distinction. As I shall be arguing that affordances and the 'general horizons' described in the quote from Husserl above are distinct, I will restrict my usage of 'horizons' to co-presented occluded surfaces and features of objects only, a restriction I shall henceforth enforce via the term 'object horizons'.)

It might be tempting to understand the horizontal structure of visual experience in terms of inferential reasoning. Such an account is suggested, for example, by Russell

(2002), who claims that we infer the shapes of objects from sense data, and Gregory (1980), who characterises perception as a form of hypothesis formation. The claim would be that we see the front of the coffee cup and infer or predict on this basis that it has a reverse side we cannot currently see. Talk of inferences and hypotheses may well be a useful metaphor for describing the neural processing implicated in perceptual cognition (see e.g. Helmholtz 1962, Fodor and Pylyshn 1981, cf. Hatfield 2002), but it fails as phenomenological description because it over-intellectualises the experience. We do not need to engage in any reasoning in order to experience the co-presence of the reverse side of a coffee cup with its facing side. In fact, if we were to engage in such reasoning, this could only be because we already possess a sense of the presence of the cup's reverse side. We might, for instance, speculate about what the reverse side actually looks like or whether the shape of the cup might differ from what we would expect given its apparent shape from our current point of view, but these questions only get off the ground given the more basic assumption that such a reverse side is there to verify or falsify our expectations. The *sheer presence* of the reverse side is built into the structure of our ordinary experience, and to describe *this* in inferential terms would lapse into the experience error. Prior to scientific or epistemological reflection on the status of the cup's reverse side, it is simply and unproblematically *there* for us. When we affirm the presence of the reverse side of the cup, though we can choose to treat the utterance as a testable prediction if, say, we are in an epistemology seminar, this is first and foremost a report about how the world presents itself in our visual experience, as Merleau-Ponty stresses

> I grasp the unseen as present, and I do not affirm that the back of the lamp
> exists in the same sense in which I say the solution to a problem exists. The

hidden side is present in its own way. It is in my vicinity. (Merleau-Ponty

1964b p.14)

If the phenomenology of the visual field outstrips what is present to the eyes at any given instant, and it is phenomenologically inadequate to explain this aspect of perceptual experience in terms of inferential reasoning on the part of the perceiver, and there are only sparse visual representations in the brain, why, then, do we experience the co-presence of occluded surfaces and unattended features of objects? According to the sensorimotor approach, the phenomenon of co-presence owes to perceivers' sensitivity to the fact that the occluded surfaces are 'available to perception through appropriate movement' (Noë 2012 p.58, italics removed), and this requires that perceivers possess what is variably referred to as 'sensorimotor understanding', 'sensorimotor knowledge' or 'sensorimotor skill'. Out of personal preference, I shall stick with the term 'sensorimotor understanding'. Sensorimotor understanding consists in an implicit grasp of what O'Regan and Noë call 'sensorimotor contingencies'. Sensorimotor contingencies are the law-like regularities which hold between sensory contents and shifts in perspective brought about through bodily movement and perturbations in one's immediate environment (O'Regan & Noë 2001 pp.940-3).

Evidence for this hypothesis comes by way of early experiments on inverted vision. Stratton (1897) constructed a pair of glasses which flipped retinal images by a hundred and eighty degrees, resulting in the visual world appearing upside down. He tested the effects of the lenses on himself by wearing them for between four and twelve hours a day over a period of eight days and kept a diary of his experiences (ibid. p.343). When not wearing the glasses, Stratton was blindfolded (ibid.). In early stages of adaptation, not only was his visuomotor coordination severely impaired, his

entire visual field was restructured and destabilised. On the first day of wearing the inverted lenses, Stratton reports: 'It did not feel as though I were visually ranging over a set of motionless objects, but the whole field of things swept and swung before my eyes' (ibid. p.344). Though initially finding himself in the grip of a 'nervous depression' (ibid. p.346) over his disjointed experience, Stratton reports a gradual honing of visuomotor control and a general stabilising of the visual field. Although the world continued to appear upside down, the experience post-adaptation is felt to be close to normal, especially when actively interacting with the world

> [T]he most harmonious experiences were obtained during active operations on the scene before me. In rapid, complicated yet practiced movements, the harmony of the localisation by sight and that by touch or motor perception - the actual identity of the positions reported in these various ways - came out with much greater force than when I sat down and passively observed the scene. (Stratton 1897 p.356)

Similarly, Kohler (1962) designed sets of goggles which distort retinal images in various ways. One set of Kohler's goggles bent the light rays entering the eye in such a way that when the head turns to the left with the eyes turned to the right, perceived objects appear to contract horizontally, and when the head turns to the right with the eyes turned to the left, perceived objects appear to expand. One of Kohler's test subjects described the initial experience of wearing these goggles as being 'as if the world were made of rubber' (ibid. p.303). After a short period of adaptation, these optical disparities completely disappear and vision is reported as returning to normal. Kohler concludes: 'Somehow the visual system has learned a general rule: a contracted image must be expanded and an expanded image must be contracted, depending on the respective position of head and eyes' (ibid.). As Kohler makes clear

75

in his introduction to his paper, he follows tradition and endorses a version of the snapshot conception of vision.[19]But having encountered compelling empirical and phenomenological reasons for rejecting the snapshot conception, the sensorimotor approach allows for a more elegant explanation: the experience is stabilised through the adaptation to a new set of sensorimotor contingencies via the honing of sensorimotor understanding; Kohler's subjects acquired a new set of perceptual skills (O'Regan & Noë 2001 pp.593-4).

This implicit grasp of sensorimotor contingencies operative in visual experience is held to be a practical, as opposed to propositional, form of knowledge (Noë 2004 pp.117-22; O'Regan and Noë 2001 p.944) ; in Ryle's idiom, it involves knowing-how (to do something) rather than knowing-that (something is the case) (see Ryle 2000 ch. II). Although we do not need to continually move in order to experience co-presence, we experience co-presence because we know implicitly how to manoeuvre our bodies in relation to the object in such a way as to bring the occluded side or feature into view

> The sensory modalities, according to [the sensorimotor approach], are constituted by distinct patters of sensorimotor contingency. Visual perception can now be understood as the activity of exploring the environment in ways mediated by knowledge of the relevant sensorimotor contingencies. And to be a visual perceiver is, thus, to be capable of exercising mastery of vision-related rules of sensorimotor contingency. (O'Regan and Noë 2001 p.943, italics removed)

---

[19] 'Vision is perhaps the most complex of the senses; nonetheless it offers the investigator a tantalising opportunity to learn how the brain processes sensory data and constructs an effective image of the outside world' (Kohler 1962 p.299).

Note that the emphasis on perceivers' *knowledge* of sensorimotor contingencies guards the sensorimotor theorist against a potential objection whereby the explanation of co-presence in terms of dispositions to act lapses into a form of behaviourism. The sensorimotor theorist does not *identify* perceptual content with behavioural dispositions, but cashes it out in terms of tacit knowledge of the sensory consequences of movements of the body and of the object in relation to the body (Jacob 2006 p.2).

Perceptual co-presence, then, is intimately bound up with our embodiment. Not only do we routinely bring parafoveal objects and occluded surfaces into view via the execution of bodily movements, it is our practical grasp of the relations between movements of the body and the objects of perception that constitutes our experience of their presence. Perhaps surprisingly, however, O'Regan and Noë have very little to say about the phenomenology of embodiment itself. This, I shall argue, generates certain troublesome but ultimately avoidable inconsistencies in their conception of sensorimotor understanding. These problems can be avoided by means of a more nuanced phenomenological analysis of the relationship between bodily agency and perceptual experience. Below I undertake such an analysis. For the remainder of this chapter I focus mainly on the work of Noë. The decision to focus henceforth primarily on Noë rather than O'Regan is motivated by the fact, noted in the introduction to this chapter, that Noë takes himself to be engaging in a form of phenomenological analysis. As such, he goes to greater lengths than O'Regan to elucidate the concept of sensorimotor understanding. Below I set out what I take to be Noë's precise commitments with regards to what sensorimotor understanding actually is and how it is exercised. I then raise three objections to Noë's account before drawing further on Merleau-Ponty's early phenomenology to sketch a more

thorough and phenomenologically satisfying interpretation of the notion of sensorimotor understanding.

6. Object horizons, affordances, and practical know-how

Noë characterises sensorimotor understanding as practical knowledge of possibilities for action. This leads him to explicate the horizonal structure of visual experience in terms of Gibsonian *affordances*. According to Gibson, to see an affordance is to directly perceive a familiar object's practical 'value' or 'meaning' (Gibson 1986 p.127), that is, to see it as suggesting a possible usage which can be taken up in action. Visual experience is laden with affordances: a chair affords sitting to a creature capable of sitting, lateral terrain affords walking to a creature capable of perambulation, the coffee cup on my desk affords grasping and drinking from, a buzzing mobile phone affords answering, and so forth. In Gibson's ecological theory of vision, although affordances are 'external' properties of objects (meaning they are properties of actual objects, rather than of mental states), they are nevertheless relational properties– they are 'animal-relative', meaning that their perceptibility depends on the behavioural repertoire of the perceiver (ibid. pp.127-8). Noë's extreme ecological proposal is that visual experience comprises affordances *all the way out*

> According to the sensorimotor view, there is a sense, then, in which *all* objects of sight (…) are affordances. To experience a property is, among other things…to experience the object as determining possibilities of and for movement. (Noë 2004. p.106, emphasis in original)

Both Noë and O'Regan tend to equivocate on the issue of whether or not these 'possible movements' need be self-initiated or not, but there is ample textual evidence to suggest that what really differentiates the sensorimotor approach to vision from more traditional cognitivist theorising is an emphasis on self-initiated action. For example, O'Regan is quick to defend the sensorimotor approach against a misreading according to which seeing always requires the exercise of a bodily action, but nevertheless states that 'action must potentially play a role' in all perception (O'Regan 2010 p.41). Noë (2012; 2010) has recently shifted away from much of the terminology of their original *Behavioural and Brain Sciences* paper, and now presents his work on perceptual co-presence under the moniker of 'actionism', the rhetoric of which is clearly indicative of an emphasis on self-initiated movement. Elsewhere he tells us, 'Only through *self*-movement can one *test* and so *learn* the relevant patterns of sensorimotor dependence' required to perceive (2004 p.13, italics in original). And it is surely only in terms of self-initiated movement that we can make sense of his otherwise bewildering comparison of visual experience, in his latest book *Varieties of Presence,* to 'a kind of dance' (Noë 2012 p.130). In any case, the invocation of Gibson's ecological theory of vision only makes sense within the context of active self-movement, so this reading is not only justified, but necessitated by the claim under consideration.

Following Husserl and Noë, I will take it as an undeniable phenomenological fact that perception is irreducibly horizonal, that is, I will grant that co-presence is a basic phenomenal feature of visual experience. What I shall criticise, however, is Noë's claim that this horizonal structure can be understood in terms of Gibsonian affordances and the possession of Rylean practical know-how. Below I present three problems incurred by Noë's characterisation of sensorimotor understanding in terms

of practical knowledge for possibilities for action. I then turn to *Phenomenology of Perception* to show how two of its major themes – the previously discussed 'body schema' and the arguably lesser known concept of 'sedimentation' – can help plug the gaps in Noë's account and form the basis of what Husserl called a 'genetic' phenomenology of sensorimotor understanding.

7. Scholastic interlude: why Merleau-Ponty?

Although Merleau-Ponty's early phenomenology is often invoked in discussions of the sensorimotor approach to perception, it is seldom the subject of sustained engagement in the sensorimotor theory literature. Appropriation of *Phenomenology of Perception* is generally limited to the occasional citation of pieces of phenomenological description with the aim of portraying Merleau-Ponty as an early advocate of the sensorimotor approach, without due attention to the wider philosophical project in which those descriptions occur (e.g. Noë 2004 p.17, O'Regan 2011 p.23). This has the doubly unfortunate consequence of portraying *Phenomenology of Perception* as a mere work of descriptive psychology and obscuring possibilities for philosophically interesting disagreement or mutual enlightenment between the sensorimotor research program and scholars of phenomenology. Hopefully, what follows sketches one path a more satisfying engagement might take by demonstrating that Merleau-Pontian phenomenology provides the conceptual resources to build upon and improve Noë's account of sensorimotor understanding and its importance to visual experience.

8. Object horizons are not affordances

There are three big problems with Noë's attempt to understand object horizons in terms of affordances for action and practical know-how. First, to perceive an affordance is to grasp a visual scene as suggesting possibilities for *voluntary* behaviour[20], and not all bodily movement that would bring co-presented features to visual presentation need be voluntary. Consider again the visual experience of a coffee cup on a desk. The anticipation that the cup has a reverse side would be fulfilled irrespectively of whether or not the movements required to see it were willed to happen. Passive, involuntary movements would do the job just as well: I could be pushed into a different spatial position or fall and land with my head on the desk behind the cup, and the perceptual anticipation would still be fulfilled. A less frivolous example would be somebody incapable of voluntary self-movement, perhaps since birth, pushed around the desk in a wheelchair while they happened to be looking at the cup. Cashing out horizons in terms of an implicit grasp of possibilities for action therefore over-specifies the content of sensorimotor understanding. Perceptual sensitivity to the way in which movements of one's body or movement of the object relative to one's body would affect one's current perceptual experience need not require seeing an object as affording possibilities for active self-movement.

With this in mind, it is important to distinguish between what we may call 'action-dependence' and 'movement-dependence' in perceptual experience. A perceptual experience is action-dependent if and only if it depends for its content or character on the perceiver knowing how to act in certain ways, that is, if it contains some reference to possible *self-initiated* movement. Affordances are clearly action-

[20] I take this to be obvious from the sorts of examples Gibson gives of affordances in *The Ecological Approach to Visual Perception*, which pertain to overt intelligent performances rather than automatic stimulus-response reflexes.

dependent in this sense, for taking up an affordance requires some exercise of the will. An experience is merely movement-dependent if it involves being sensitive to sensorimotor contingencies without associating those contingencies with possibilities for self-initiated movement. Horizons are certainly movement-dependent, but they are not thereby action-dependent. Hence, there is some conceptual space between affordances –which presuppose volition, and object horizons, which do not, which gives us good grounds for refraining from equating the two. Note that this is not to say that object horizons are not explicable in sensorimotor terms. The sensorimotor theorist can continue to hold that it is necessary for experiencing horizons that one's visual system be sensitive to movement of one's body and nearby objects and/or to movement of objects relative to the perspective imposed by one's body. The upshot is simply that this does not by itself license Noë's far stronger claim that object horizons *are* affordances.

A second problem with equating object horizons with affordances concerns the psychological development of vision and agency. There are empirical grounds for holding that object horizons are developmentally prior to affordances. Consider Held and Hein's famous 'kitten carousel' experiment (Held & Hein 1963). In contrast to a popular misconception the experiment concerned visually guided behaviour, not visual experience per se. For three hours daily ten pairs of neonatal kittens were placed in apparatus resembling a fairground carousel: a circular box with a two-pronged rotating arm at the centre. At one end of the rotating arm, the 'active' kitten of the pair was attached to a harness with its feet in contact with the ground so that it could control its own locomotion. The 'passive' kitten was placed in a box with its head held in a fixed position and suspended from the other end of the rotating arm so that it could see the inside of the apparatus but not move around freely (although it

could still move its own eyes). By walking, the active kitten pulled the passive kitten around the carousel, so while both sets of kittens were exposed to the same patterns of movement and visual stimuli, only the active kittens' movements around the apparatus were self-initiated. Held and Hein found that the active kittens avoided visual cliffs, put out their paws to brace themselves when picked up and placed on a different surface, and displayed avoidance behaviour when faced with looming objects, while the passive kittens did not, though their responses did normalise within forty-eight hours (ibid. p.875). Unsurprisingly, they conclude that 'self-produced movement with its occurrent visual feedback is necessary for the development of visually guided behaviour' (ibid).

The passive kittens' normal pupillary reflexes, healthy eyes, and the quickness with which they adapted to visually guided behaviour indicate their visual sense was not impaired by lack of self-movement; rather their ability to coordinate voluntary movement with their visual experience was temporarily hindered (ibid. p.875-6). Noë interprets the passive kittens' failure of the visual cliff task as evidence for a lack of depth perception (Noë 2004 p. 234 §9). But given their intact visual system this seems implausible, and such an interpretation incurs the potentially intractable, perhaps even paradoxical, problem of explaining how a creature could see three-dimensional objects in their immediate environment without experiencing depth. This would be tantamount to asserting that despite their fully-functioning visual systems, the passive kittens see the cliff in two-dimensions – an ad hoc stipulation if ever there was one. A more conservative explanation, suggested by Kinsbourne, is that the cliff looks the same to both kittens, but only for the active kittens has the appearance of the cliff come to be integrated with the feel of solid terrain under their paws (Kinsbourne 1995 pp.215-6).

Continuing to grant that object horizons are a basic, irreducible feature of visual experience, engaging in a spot of feline 'hetrophenomenology'[21] allows us to draw the following moral: the passive kittens perceived objects as objects, and therefore experienced horizons, but, unlike the active kittens, could not perceive affordances; visual cliffs and looming objects did not 'negatively afford' avoidance (see Gibson 1986 p.137), the approaching floor did not afford paw-extension, etc. The ability to see affordances developed as their spatial vision and capacities for bodily action were allowed to integrate as they otherwise would have naturally. Hence, there is good reason for thinking that movement-dependent object horizons are phenomenologically more basic than action-dependent affordances as the former can apparently exist in the absence of the latter but not vice versa, and also that the ability to see affordances requires some additional development over and above a more primitive capacity to experience object horizons. Again, this is not to say that the passive kittens' visual experience cannot be understood in sensorimotor terms. They were, after all, exposed to the same patterns of movement-dependence as the active kittens, only their movement was almost entirely involuntary. Movement dependence is still a bodily phenomenon, as it consists in perceptual sensitivity to the effects of (passive) bodily movements on visual content and movement of objects relative to an embodied perspective. The point is that mere visual sensitivity to movement does not equate to an ability to grasp affordances for action.

A third problem with equating object horizons with affordances stems from a latent contradiction in the attribution of practical knowledge to subjects incapable of performing the required bodily movements for themselves. While Noë does not

---

[21] The attempt to sketch a rational reconstruction of the life-world of something (a person, animal, machine, etc) by interpreting observations of its behaviour in terms of our own conceptual scheme (see Dennett (1991) pp.72-85).

claim that severe restrictions on a perceiver's ability to act would result in blindness (which would be patently and demonstrably false), he does claim that the preservation of normal vision in the paralysed owes to the retention of sensorimotor understanding

> Paralysis is certainly not a form of blindness…Even the paralysed, whose range of movement is restricted, understand, implicitly and practically, the significance of movement for stimulation. They understand, no less than those who are not disabled, that movement of the eyes to the left produces rightward movement across the visual field, and so forth. Paralysed people can't do as much as people who are not paralysed, but they can do a great deal; whatever the scope of their limitations, they draw on a wealth of sensorimotor skill that informs and enables them to perceive. (Noë 2004 p.12)

Recall that Noë is wont to describe sensorimotor understanding as a form of non-propositional practical knowledge, or skill (ibid. pp.117-22). On a standard conception of practical knowledge or skill, knowing-how to ϕ necessitates being able to ϕ. This conception of practical knowledge is certainly what Ryle had in mind in his original articulation of his knowing-how/knowing-that distinction in *The Concept of Mind*, as he argued that skills are acquired dispositions to act (Ryle 2000 p.33). For example, if I cannot play the guitar to a certain standard, then I do not yet possess the skill of guitar playing: I do not know *how* to play the guitar in the required sense. If in time I learn to play to an adequate standard, but then an injury requiring physical rehabilitation renders me unable to execute the required movements, I lose my *practical* knowledge, even if I can describe quite well what it is I am supposed to do with the instrument to produce the desired sounds. Of course, there is a spectrum of cases here. I might struggle to play a certain riff if I have

recently trapped one of my fingers in a car door, but it would be premature to conclude that I have lost my musical skill. Assuming I make a full recovery, I might simply wake up one morning having had time to heal, and pick up the guitar and play the riff without difficulty and without the need for any tedious relearning process. In this instance, I have retained my skill - an acquired disposition to act in a particular way - but the circumstances required for the actualisation of this disposition, a necessary part of which are unbroken fingers, were prevented from coming to be. But there has to be a pragmatic cut-off point with this kind of reasoning. The longer a skill goes unexercised and the more elaborate and unlikely the circumstances required for the skill to be put to use become, the less plausible it is to continue to ascribe the skill to the person in question. If, for example, the circumstances required for a skill to be applied include something as extreme as the regrowth of a dismembered limb, the continued ascription of the skill to the unfortunate subject is fantastical. Though self-regenerating arms and legs might well be metaphysically or even physically possible, the sheer improbability of these circumstances coming to be, necessary as they are for the skilled performance in question, renders the continued ascription of the skill borderline nonsensical. Skills are, in Merleau-Ponty's words, 'knowledge in the hands' (Merleau-Ponty 2012 p. 145), and this is why we do not say of an athlete past her prime that she has retained her skill even though she can no longer compete, but rather that she can no longer compete because, alas, age has deprived her of her skill. Even if this Rylean characterisation of skill is not to the reader's satisfaction this need not matter because, as the following quotation makes clear, it is explicitly endorsed by Noë[22]

---

[22]Ryle's knowing-how/knowing-that distinction is challenged by Stanley and Williamson (2001), but given Noë's endorsement of the distinction I shall grant it for the sake of argument. My intuition on this matter is that Stanley and Williamson's

I would have thought that if a ski instructor can't do the jump, then she

doesn't know how to do it...She knows how the jump is done, but not how to

do it. Sadly the same is true of the pianist [who has lost an arm]. He may

retain all sorts of cognate [propositional] knowledge (…) but when he lost his

arms, he lost his know-how. For the knowledge was, precisely, arm-

dependent. (Noë 2004. p.121)

The problem for Noë's sensorimotor theorist should now be obvious: if the

possession of practical knowledge is necessarily dependent upon or identical to an

ability to act, then it is nonsensical to attribute practical knowledge of possibilities

for action to those in whom such abilities are lacking.[23]

Locked-in syndrome is instructive in this regard. The most common variation of

locked-in-syndrome, the so-called 'classical' variation, involves complete paralysis

apart from blinking and limited vertical eye movement (Bauer, Gersenbrand and

Rumpl 1979)[24]. Sufferers of the syndrome can communicate using systems of blinks

and vertical eye movements and with the help of various eye-tracking technologies

---

critique, which concerns the logical form of knowledge ascription sentences, rather
misses the point of Ryle's original distinction, which is more phenomenological than
logical. Ryle himself favours such an interpretation when he writes that *The Concept
of Mind* 'could be described as a sustained essay in phenomenology, if you are at
home with that label' (Ryle 2009 p.196).

[23] This line of criticism can applies equally to the position of Schellenberg (2007),
who argues that spatial perception is dependent upon perceivers' capacity to act.

[24] Classical locked-in syndrome differs from 'incomplete' and 'complete' variations
of the condition. With incomplete locked-in syndrome, a very small amount of
additional motor control is preserved, while paralysis in the complete variation
extends even to blinking and vertical eye movement (Bauer, Gersenbrand and Rumpl
1979). Sufferers of complete locked-in syndrome are unable to communicate, so little
is known about their experience in comparison to the other two variations, although
the familiar phenomenon of sensory fatigue, where continual presentation of
identical stimuli results in a decrease in the corresponding sensation, suggests their
vision may deteriorate.

(Laureys et al 2005), so rather a lot is known about their experience from firsthand reports. People with locked-in syndrome retain full visual consciousness and their intellectual capacities remain untouched. Indeed the condition's defining characteristic is the patient's being 'literally locked inside his body, aware of his environment but with a severely limited ability to interact with it' (Patterson & Grabois 1986 p.758). Although locked in syndrome can sometimes negatively affect visual attention (Smith & Delargy 2005 p.406), one patient goes so far as to describe his vision as 'normal, if not enhanced' (Chisholm & Gillet 2005 p.94). As they see perfectly well we can say without controversy that locked-in perceivers' experience has a horizonal structure. Of course, sensorimotor theorists need not, do not, and given their phenomenological starting point cannot, deny this. But it is wrongheaded for Noë's sensorimotor theorist to explain the locked-in subject's visual experience in terms of practical knowledge of how 'movement of the eyes to the left produces rightward movement across the visual field' given that, on Noë's own account, their inability to perform these movements renders incoherent the attribution of the practical knowledge required to make them. This goes a fortiori for more complex interactions. The locked-in patient does not know *how* to manoeuvre their body around an object in the sense of having the required skills; were they miraculously cured it is likely that they would need to reacquire them through rehearsal and physiotherapy. Practical knowledge - at least as Noë understands it - is therefore just the wrong sort of thing to account for object horizons.

At this point, the following question becomes pertinent: If seeing an affordance requires practical know-how, what are we to say of the perception of affordances for the locked-in perceiver? It would be wildly counterintuitive to suggest that locked-in syndrome patients, who have lost almost all their practical knowledge, thereby

cannot see affordances. Their visual experience is not that of a human equivalent of Hein and Held's passive kittens. Having been accustomed to living a life of practicality, and given the full preservation of their intellectual and visual capacities, it would be ad hoc and implausibly farfetched, not to mention rather offensive, to attribute to them an impoverished consciousness whereby they no longer see chairs as for sitting, doors as for opening and closing, coffee cups as for filling and drinking from, etc. We must therefore reject not only the conflation of horizons with affordances, but also the implied conflation of the capacity to see affordances with the possession of Rylean practical knowledge. What is needed is a better philosophical framework in which to make sense of sensorimotor understanding. Merleau-Ponty's early phenomenology gives us the conceptual tools to do just this.

9. Motor signification, sedimentation and the body schema

These are the facts to be accounted for: object horizons are not Gibsonian affordances, but through the garnering of practical knowledge, they may be 'upgraded' (so to speak) to affordances. But the perceived practical value which the objects of vision have for the perceiver who possesses and can exercise the relevant know-how – affordances – persists even after this know-how has been lost. The sensorimotor approach is therefore hampered by Noë's restricted conceptual toolkit. Fortunately, Merleau-Ponty's *Phenomenology of Perception* gives us the additional tools needed to untangle the knots in which Noë's sensorimotor theorist ties herself by attempting to conceive of sensorimotor understanding solely in terms of Rylean practical know-how and Gibsonian affordances, though understanding precisely what Merleau-Pontian phenomenology has to offer the sensorimotor theorist calls for

some preliminary exposition on the aim of phenomenology considered not as a subject matter for philosophers of mind and cognitive scientists, but as a discipline.

Phenomenological philosophy is personal-level analysis par excellence, but, as all phenomenologists know, there is considerably more to it than introspective reports on the content or character of one's own psychological states. Husserl's philosophical project gradually evolved from what he termed 'static' into 'genetic' phenomenology, and Merleau-Ponty's *Phenomenology of Perception* sits squarely in the latter category. Static phenomenology describes experience in terms of hypothetically hypostatised appearances, arguably much like the personal level version of the snapshot conception of vision. Genetic phenomenology goes a step further and escapes the shortcomings of the snapshot conception, by attempting to trace the origins of these appearances in lived experience. Alternatively put, we can say that while genetic phenomenology aims to uncover the structures of consciousness through which appearances are formed or, in phenomenological jargon, 'constituted', static phenomenology is (at best) limited to the description of its end products. Interestingly, Noë, who takes himself to be 'investigating the phenomenology of perceptual experience' (Noë 2004 p.33), echoes Husserl's move from static to genetic phenomenology when he says 'the task of phenomenology ought to be not so much to depict or represent or describe experience, but rather to catch experience in the act of making the world available' (ibid. p.176). It is fitting, therefore, that the shortcomings of Noë's account of sensorimotor understanding can be rectified by adopting a genetic-phenomenological approach, that is, by going beyond the description of the content of perceptual states in order to uncover the subjective operations through which such states come to be constituted in lived experience. Two Merleau-Pontian concepts are indispensable in this regard: the

already discussed body schema, and a phenomenon Merleau-Ponty calls 'sedimentation'.

Recall that in contrast to a percept or mental representation of one's own body (a body *image*), the body *schema* is Merleau-Ponty's term for the integrated system of pre-reflective motor capacities which structure lived experience. Though Merleau-Ponty would certainly agree with Gallagher (2005) that the body schema should be distinguished from any mental state which has the body as its object[25], the body schema has for Merleau-Ponty a genetic-phenomenological significance over and above its being a 'system of sensory-motor processes that constantly regulate posture and movement that function without reflective awareness or the necessity of perceptual monitoring' (ibid. pp.37-8). As Merleau-Ponty puts it, the body schema is not itself an appearance or an object of thought, but a 'law of constitution' (Merleau-Ponty 2012 p.101), meaning that it conditions the ways in which things appear to the perceiver. The body schema is therefore indispensable to a genetic-phenomenological analysis of embodied perceptual experience. Recall Gibson's characterisation of affordances as 'animal relative'. The notion of a body schema further illuminates this point. It is by virtue of having a body schema that objects can afford usage and one's environment can take on a practical significance, because the way in which a subject can interact with their environment is relative to the range of possible actions permitted by their particular bodily morphology. A body schema is therefore a precondition of the formation of affordances. This is the meaning of Merleau-Ponty's remark that 'my own body is the primordial habit, the one that conditions all others and by which they can be understood' (ibid. p.93).

---

[25] And again: 'A body image consists in a system of perceptions, attitudes, and beliefs pertaining to one's own body. In contrast, a body schema is a system of sensory-motor capacities that function without awareness or the necessity of perceptual monitoring.' (Gallagher 2005 p.24)

Crucially, the body schema is adaptable. By honing skills and acquiring new habits, it can be 'reworked and renewed' (ibid. p.143) (I discuss the adaptability of the body schema in greater depth in the next chapter). Perceivers' capacities to perceive the world in new ways, which are informed by their acquired skills and other changes in bodily morphology, endows their experience with a uniquely bodily kind of *meaning* which Merleau-Ponty calls 'motor signification' (ibid p. 113). Though it is rather tricky, particularly for an analytically trained philosopher, to define or articulate precisely the idea of 'bodily meaning', an example of the body schema's contribution to lived experience should help to clarify what Merleau-Ponty has in mind. Though saying so perhaps runs the risk of sounding like a musical chauvinist, the perceptual encounters of expert musicians with their instrument of choice are significant in a way in which the non-players' perception of the same instrument are not. For example, a skilled guitarist has at her disposal a certain 'muscle memory' of (inter alia) various chord formations, picking techniques and scale shapes, as well as a working knowledge of what constitutes a slim or thick neck, a solid or weight-relieved body, heavy or light gangue strings, or an awkward or comfortable (relative to their preferences and playing style) layout of volume and tone pots, tremolo systems, and even tuning pegs, which are lacking in non-musicians who are nevertheless perfectly capable of performing identical finger movements. Consequently, guitars are perceived in more complex and meaningful ways by guitarists – they draw on a richer sensorimotor understanding alien to non-players. The non-guitarist knows (in the propositional sense of 'knows' – knowing-that) very well what guitars are for, and might even know something of how they are played and what the functions of their various pieces of hardware are, but a guitar does not afford playing (or other meaningful interactions such as tuning, lowering bridge

height or adjusting the truss rod) for them in the concrete sense experienced by the skilled player for whom the guitar represents a manifold of genuinely possible motor projects. Unlike non-players, the seasoned player has with respect to guitars what Merleau-Ponty calls the 'power to reckon with the possible' - they can orient themselves in relation to various dealings with guitars which for them, by virtue of their refined understanding of the instrument, are live possibilities (ibid. p.112) (see Romdenh-Romluc 2007; 2011b pp.62-192). Similar morals apply, mutatis mutandis, for other forms of skilful sensorimotor interactions. As a non-driver, the interior of a car is alien and intimidating to me in a way which invites laughter from my road-able friends even though I have ridden shotgun countless times, and a recent trip abroad served as a lesson in how, despite comprehending and obeying the instructions of helpful local residents, my lack of familiarity with foreign methods of public transport amounted to a kind of behavioural illiteracy. Doubtless we can all recall similar experiences where one struggles to 'interpret' one's environment while others negotiate it effortlessly. The difference lies not between two different bodily morphologies, but between the manner and degree to which the same surroundings call for different kinds of bodily engagement. It is by virtue of the body schema and its adaptability that we are 'geared' into our environment (as Merleau-Ponty, paraphrasing Heidegger, sometimes puts it) in such a way that it *makes practical sense* to us. Merleau-Ponty's notion of the body schema therefore provides insight into the phenomenological constitution of Gibsonian affordances, as well as helping to make sense of Gibson's own construal of affordances as both properties of external objects and visible 'values or meanings' (Gibson 1986 p.127).

With this in mind, we can turn to our second key Merleau-Pontian concept. Normal (i.e. typically developed and non-pathological) subjects can integrate prior mental

operations into their behaviour in such a way as to alleviate the need for any rehearsal of the reasoning behind them. Merleau-Ponty's term for this phenomenon is 'sedimentation'

> These acquired worlds which give my experience its secondary sense, are themselves cut out of a primordial world which grounds the primary sense of my experience. Similarly there is a "world of thoughts", a sedimentation of our mental operations, which allows us to count on our acquired concept and judgements, just as we count upon the things that are there and that are given as a whole, without our having to repeat their synthesis at each moment. (Merleau-Ponty 2012 p.131)

Merleau-Ponty's discussion of language provides a handy illustration of sedimentation at work in everyday experience (see ibid. pp.179-205). Learning the correct way to use a word, coining a phrase, or adopting a manner of talking in line with a newly acquired attitude (such as moderating one's language for the sake of political correctness) are initially the outcome of a creative thought processes through which new habits are constituted. Merleau-Ponty calls this spontaneous or expressive use of language, 'speaking speech'. With practice, the speaker comes to utilise the new linguistic device without needing to remind themselves of the reasoning behind their wording, sometimes having even forgotten it entirely. We can forget the origin of a phrase or the process through which we learned to use it while continuing to routinely deploy it correctly. Merleau-Ponty calls this sedimented (sic) linguistic usage, 'spoken speech'. Merleau-Ponty sometimes presents spoken speech as a derivative, secondary and therefore inferior or 'inauthentic', form of linguistic communication (or at least he does so in *Phenomenology of Perception,* cf. Merleau-Ponty 1964c), but we need not accept this, as the two are mutually grounding.

Though sedimented spoken speech is born of speaking speech, spontaneous speaking speech cannot occur ex nihlo, as a novel linguistic coinage presupposes an extant set of word meanings and connotations to be modified and re-appropriated, and against which the novel contribution of a new linguistic creation can be understood (see Baldwin 2007 for insightful criticism of Merleau-Ponty along these lines). Hence, building up the layers of meaning through which subjects engage with the world and each other, rests on a 'double moment of sedimentation and spontaneity' (Merleau-Ponty 2012 p.132) – the formation of new modes of self-expression and communication within the confines of established linguistic practices.

Returning now to the three problems with Noë's account of sensorimotor understanding outlined in the previous section, and bearing in mind the discussion of the body schema above, my modest suggestion is that similar Merleau-Pontian morals apply to motor significations as to linguistic meaning. Once a perceiver has acquired a piece of practical knowledge – a skill – through an adaptation of their body schema, the perceived world gains for them a new motor signification and the perception of a novel affordance is made possible. However, just as we continue to use and understand 'spoken speech' without recollection of the creative 'speaking speech' through which it was constituted, we can continue to grasp the motor signification of a familiar object once the skilful know-how from which it originates has been lost through injury, illness, or in the absence of the enabling condition(s) required in order for it to be exercised successfully.

With this genetic phenomenological framework in place, we are now poised to supplement Noë's account of sensorimotor understanding and deal with my three objections. Contra Noë, object horizons are not themselves affordances for action. Keeping Held and Hein's passive kittens as our example, we may say that the ability

to perceive affordances – to perceive objects in one's environment as exhibiting a motor signification– is the product of the development of practical know-how via adaptations of one's body schema. This is why despite already being able to see object horizons, the passive kittens did not grasp affordances until they had honed the relevant bodily skills. Hence, the perceptual meaningfulness of affordances is constituted (in the genetic-phenomenological sense of 'constitution' –as coming to appear as such) through skilful sensorimotor interactions, of which the body schema is the vehicle. Now recall the locked-in syndrome patient, for whom the practical knowledge or skill required to take up an affordance is lost. Their meaningful relationship with their visual world – their system of motor significations - is not lost, despite their deficit, because for them affordances are already constituted – the bodily meaning of their familiar environment is *sedimented* (sic) – although the opportunity to form new motor significations is now largely closed to them due to their severely restricted possibilities for novel sensorimotor interactions through which new motor significations could be constituted.

In closing this section, two additional passages from *Phenomenology of Perception* will serve to further illustrate the multi-level conception of meaningful sensorimotor understanding discernable in Merleau-Ponty's work that is lacking in Noë's. The first recalls Merleau-Ponty's notion of the 'intentional arc' from his first book, *The Structure of Behaviour* (Merleau-Ponty 1967). The second comes from his discussion of (non-aplastic) phantom limbs

> [T]he life of consciousness – epistemic life, the life of desire, or perceptual
> life – is underpinned by an "intentional arc" that projects around us our past,
> our future, our human milieu, our physical situation, our ideological situation,

and our moral situation, or rather, that ensures that we are situated within all these relationships.' (Merleau-Ponty 2012. p.137)

What refuses the mutilation or the deficiency in us is an I that is engaged in a certain physical and inter-human world, an I that continues to tend toward its world despite deficiencies or amputations and that to this extent does not *de jure* recognise them. The refusal of the deficiency is but the reverse side of our inherence in a world, the implicit negation of what runs counter to natural the movement that throws us into our tasks…to have a phantom limb is to remain open to all of the actions of which the arm alone is capable and to stay within the practical field one had prior to the mutilation. (ibid. pp.83-4)

The locked-in patient's sensorimotor understanding, though initially the product of practical knowledge, no longer depends for its continued existence on bodily skills as such, but rather inheres in their visual experience as a sedimented 'projection' of value or, to use a less extravagant phrase, an established way of seeing informed by past bodily engagements with the world. And just as the amputated arm survives for the amputee as a phantom so long as they continue to live through their familiar situation, with all the established affordances they have built up through the skilful use of their now-absent limb, the locked-in perceiver continues to 'project around them their past…human milieu…and physical situation' and thereby preserve the meaningful structure of their perceptual experience. There is therefore what might be described as an 'historical' dimension to the phenomenon of sensorimotor understanding which cannot be adequately captured by the language of commonsense psychology and ecological optics, to which Merleau-Ponty's phenomenology gives voice. And in so doing, it dissolves the worries incurred by Noë's undercooked phenomenology. While Noë is certainly correct that 'the task of

phenomenology ought to be…to catch experience in the act of making the world available' (Noë 2004 p.176), this requires explicating not just how perceivers 'bring the world forth' (Noë 2012 p.14) by actively applying their sensorimotor understanding in experience, but also what it is to be *passively* embodied and situated in such a way as to make such understanding possible in the first place.

This conclusion echoes Husserl's distinction between active and passive genetic-phenomenological constitution. As Husserl's project evolved from static into genetic phenomenology, he became increasingly sensitive to the fact that phenomenological constitution encompasses both perceptual acts performed by the perceiving subject and a background of passive openness to the world against which such acts can operate, and that a thorough analysis of the activity of perceptual consciousness eventually leads us back to this more basic, passive level of experience

> [A]nything built by activity necessarily presupposes, as the lowest level, a passivity that gives something beforehand; and, when we trace anything built actively, we run into constitution by passive generation. (Husserl 1988. p.78)[26]

In light of this, I think that the unhelpful mantra of 'perception is something we do', rampant in the sensorimotor theory literature and the subject of much derision among detractors, deserves to be jettisoned, but at no real cost to the sensorimotor

---

[26] Crucially, Husserl is *not* here espousing a form of classical empiricism whereby perception consists in the passive reception of impressions or sense-data (Hume 1977 §2; Russell 1912 ch. I) (though he may well do elsewhere). This would fall foul of Sellars' (1997) 'myth of the given'. Husserl's distinction between activity and passivity is not between pure sensation and conceptual thought, but between a *pre-reflective* openness to the world and acts of *reflective* consciousness. The issue of conceptual versus non-conceptual content, upon which I take no stance, is orthogonal to this less contentious point.

approach.[27]Noë in particular seems to assume that either perceptual experience must be conceived as wholly active – as skilful interaction with one's environment, or else as wholly passive - as the construction of internal pictures of the world out of afferent sensory data a la the snapshot conception (Noë 2004 ch. II). But this is a false dichotomy. The fact that the richness of our visual experience owes to the possession and exercise of sensorimotor understanding does not entail that perception *just is* a form of action, and allowing for passivity in experience does not entail lapsing back into the snapshot conception. Nor does the admission passivity reduce the role of embodiment in structuring perceptual experience. If anything, accommodating something like the Husserlian distinction between active and passive constitution, as I believe Merleau-Ponty does in his phenomenology of embodiment (see e.g. Morris 2010), allows us to make better sense of the intimate relationship between bodily agency and perceptual experience, without the distractions of radical rhetoric (hence my reluctance, mentioned briefly at the outset, over the label of 'enactivism'). The sensorimotor theorist who, like Noë, also aspires to be a phenomenologist has their work cut out for them, but it is hoped that the above considerations might be a step in the right direction.

10. 'Sensorimotor chauvinism?'

---

[27] O'Regan and Noë's unflinchingly reiterate this claim. They declare 'experiences…are not states. They are ways of acting. They are things we do.' (O'Regan & Noë 2001 p.960' and later, 'Visual consciousness is not a special kind of brain state, or a special quality of informational states in the brain. It is something we do' (ibid. p. 970). Noë's *Action in Perception* opens with it, 'The main idea of this book is that perceiving is a way of acting. Perception is not something that happens to us, or in us. It is something we do' (Noë 2004 p.1) while O'Regan's *Why Red Doesn't Sound Like a Bell,* tells us that perception's qualitative character or 'feel' 'is something we do' (O'Regan 2011 p.68)

Thus far I have appealed to a congenial mixture of phenomenology and cognitive science to argue that the structure of perceptual experience owes to an implicit understanding of movement-dependent sensorimotor contingencies and action-dependent motor significations which are constituted through bodily engagements with the world. In this section I defend the general underlying claim of this chapter - that perceptual experience is structured by our embodiment - against Andy Clark's accusation of 'sensorimotor chauvinism'. Clark argues that the sensorimotor approach risks overselling the contribution of embodiment to perceptual experience (2008 ch.VIII; 2006; 2001; Clark & Toribio 2001). Though some of Clark's objections focus on certain specific and non-essential aspects of the sensorimotor approach I have not endorsed, he does raise a more general worry about the extent to which visual experience is determined by sensorimotor understanding. Responding to this will necessitate a shift of emphasis away from phenomenological analysis towards more straightforwardly empirical lines of argumentation.

Clark draws on well-known empirical work by Milner and Goodale in support of their 'dual-stream' hypothesis (Milner & Goodale 1995; 1992) (see also Jacob & Jeannerod 2003). Milner and Goodale synthesise a wealth of behavioural and neurological evidence to support the hypothesis that visual information in the brain is divided between two distinct streams of processing. The *ventral* visual stream processes information pertaining to object location and classification, and is therefore often glossed as the 'vision-for-perception' stream. The *dorsal* visual stream processes visual information relevant to the guidance of action, and is often glossed accordingly as the 'vision-for-action' stream. In certain neurological conditions, the interaction between the two streams is disrupted, resulting in some unusual pathologies. One of Milner and Goodale's patients, referred to simply as 'DF',

suffered damage to the part of her brain housing the ventral stream as a result of carbon monoxide poisoning, leaving her with visual form agnosia. DF is unable to visually identify or locate objects, but she can recognise objects by touch, accurately grasp them on command, and even perform well on certain visuomotor tasks such as posting her hand or a piece of card though a letterbox-shaped hole (Milner and Goodale 1995 p.129). On a related note, there exist more extreme cases where subjects can engage in all manner of visually-guided behaviours, including walking around and even making a cup of tea, in the total absence of visual awareness and without any recollection of having done so, as occurs in instances of epileptic automatism brought on by frontal lobe seizures (Damasio 2000 pp.96-101). Conversely, people suffering from optic ataxia caused by lesions to the area of the cortex where the dorsal stream is located can visually identify and locate objects very well but struggle to use the information garnered from conscious vision to guide their actions. Optic ataxia patients cannot, for instance, correctly orient their hand or a piece of card in relation to the same letterbox-style slot, despite being able to make accurate verbal reports as to the shape and location of the target on the basis of their visual experience (Milner and Goodale 1995. p.98).

Clark also cites an experiment conducted by Aglioti, Goodale and DeSouza (1995), this time on non-pathological subjects, using a three-dimensional version of the Ebbinghaus (or Titchener circles) illusion. The Ebbinghaus illusion occurs when subjects are presented with two images of small circles of the exact same size surrounded by rings made up of other circles. The surrounding circles in one image are notably larger than those in the other. Despite the middle circles being the same size in each image, the one surrounded by the ring of smaller circles appears larger (for an optical explanation of the Ebbinghau illusion, see Roberts, Harris and Yates

2005). Aglioti, Goodale and DeSouza constructed a three dimensional model of the illusion out of differently sized poker chips and instructed subjects to pick up the chip in the middle. They report that their test subjects' grasping of the centre poker chips was accurate despite their being subject to the visual illusion, suggesting that the processing involved in the performance of this visuomotor task is restricted to the dorsal stream and is not influenced by the apparent shape of the chip as it appears in their visual experience, which is presumably processed in the ventral stream. If it is true that conscious vision and visually-guided action are the products of two completely separate cognitive subsystems then it could look like bodily action and visual experience are far more easily divorceable than the sensorimotor theorist would like. Clark therefore follows Miler and Goodale in holding that visually-guided action and visual experience are doubly dissociated; it is the ventral visual stream which is responsible for visual consciousness, while the dorsal visual stream operates unconsciously to guide action

> [W]hat is at issue is not just the evidence of substantial dissociation [between perception and action] but the best functional and architectural explanation of the evidence. And the best functional and architectural explanation, according to Milner and Goodale, and others, is that conscious perceptual experience reflects the activation of representations that have less to do with the fine details of world-engaging sensorimotor loops and more with the need to assign inputs to categories, types and relative locations so as to better sift, sort, select, identify, compare, recall, imagine and reason. (Clark 2008 p.192, italics removed)

In passing, Noë remarks that the dual streams hypothesis is 'at best orthogonal' to his version of the sensorimotor approach to visual experience (Noë 2004 p.19). I think

that Noë is correct but that he fails to spell out precisely why this is so, so below I attempt to clarify this on behalf of the sensorimotor theorist.

First of all, the fact that sensorimotor understanding, be it of the kind implicated in the task of posting a letter through a letter box or in walking into a kitchen and making a hot drink, can be present and operative in the guidance of action in the absence of visual awareness does not mean that the same sensorimotor understanding does not profoundly shape visual consciousness *when it is there*, so visual form agnosia and epileptic automatisms, where ex hypothesi there is no visual awareness, prove nothing with respect to the role of sensorimotor understanding in determining the content of ordinary visual experience.[28] The onus therefore lies with Clark's appeal to Aglioti, Goodale and DeSouza's three-dimensional Emmingbaus illusion and to the alleged functional dissociation between perception and action in optic ataxia. Fortunately for the sensorimotor theorist, neither of these suffices to discredit the claim that visual experience is dependent for its content on sensorimotor understanding.

With regards to Clark's appeal to the three dimensional Emmingbaus illusion, there are reasons for doubting the extent to which Aglioti, Goodale and DeSouza's data hold up to empirical scrutiny. Franz and colleagues (V. Franz, Gegenfurtner and Scharnowski 2005; V. Franz 2003) undertook a similar experiment using an aluminium disk as the centre circle, which was placed in the middle of a drawing of

---

[28] Note that this does not necessarily cut against O'Regan's solution to the hard problem either. For O'Regan, 'sensorimotor contingencies determine the quality of experience' (O'Regan 2012 p.24, italics removed), while awareness of the quality is explained via a 'hierarchical form of cognitive access similar to that used in "higher-order" theories [...] of consciousness' (ibid. p.25). Though the relevant sensorimotor contingencies are present in cases of epileptic automatism and visual agnosia, it remains open for O'Regan to argue, quite plausibly, that it is the lack of the relevant form of cognitive access, caused by the patients' seizure or brain damage, which explains the absence of normal visual awareness.

the surrounding circles on a piece of cardboard. Rather than relying on observing

video footage, they used special sensors fitted to the subjects' hand to track their

finger movements in order to ensure that the illusion effect was not 'contaminated' by

additional grasping movements made immediately after they had touched the disk

(V. Franz, Genenfurtner & Scharnowski 2005 p.1361). They found that the effect of

the illusion on grasping was *not* corrected for until *after* the disc was touched. This

was found to be the case both when subjects could see their own hands and when the

apparatus was sent up in such a way that they could not (ibid. p.1374). As they

themselves point out, this directly contradicts Aglioti, Goodale and DeSouza's

reports (ibid. p.1376-7). Hence, Milner and Goodale's assertion that 'grip size [is]

determined by the true size of the target' and is 'uninfluenced by the illusion' (Milner

& Goodale 1995 p.168) and Clark's subsequent conjecture that 'the processing

underlying visual awareness may be operating quite independently of that underlying

the visual control of action' (Clark 2008 p.185) is at best empirically questionable, at

worst empirically false. It would therefore be unwise to hang any substantial

philosophical point about the severability of perceptual experience and sensorimotor

understanding on this particular empirical example.

With regards to Clark's appeal to Milner and Goodale's take on optic ataxia, it is far

from clear that it can do the philosophical work Clark requires of it. What Clark

would need is a reason for thinking that optic ataxics enjoy rich visual experiences

despite a complete lack of sensorimotor understanding. But the fist of these is only

partially true, and the second is patently false. Optic ataxics, who have dorsal stream

lesions, *do* in fact exhibit perceptual deficits, as the sensorimotor theory would

predict, namely in peripheral vision (that is, their peripheral visual awareness is even

worse than the rest of us!), which negatively affects their visual attention (Pisella et

al 2009 p.3040; 3042). In O'Regan's terminology, their vision is less 'grabby' than that of ordinary perceivers. And, as Rossetti, Pisella and Vighetto (2003) note in a detailed review of the clinical literature on optic ataxia, not only can optic ataxics perform a great variety of visually-guided actions without difficulty, many of them tend not to suffer failures of visuomotor control in their daily lives. When they do, the deficits tend to be specific to particular tasks and are usually only briefly noticed (bid. p.176). Furthermore, optic ataxics' coordination improves considerably when they are allowed to take more time to complete a visuomotor task, suggesting that the effects of optic ataxia should be understood more conservatively, in terms of a breakdown in motor coordination on the basis of visual information, rather than any strict dissociation between the content of their visual experience and their capacities for visually-guided action

> Optic ataxia patients can perform all types of natural actions and are impaired only when a time constraint is imposed: their on-line motor updating is altered and their reaches toward peripheral targets are impaired when they have to produce immediate responses toward unknown objects... optic ataxia does not appear as a general deficit of action but rather as a specific deficit localized at a restricted level of action organization that is immediate visuomotor control. (ibid. p.177)

So, with regards to optic ataxia, the facts are these: optic ataxics' perceptual experience is fairly normal, but not quite identical to that of non ataxics, and their visuomotor coordination is notably off when temporal constraints are imposed on a visuomotor task. But this has absolutely no bearing on the claim that their perceptual experience is structured by a rich sensorimotor understanding comprising movement-dependent object horizons (which they see) and action-dependent motor

significations or affordances (upon which they can act). We should therefore remain guarded against any attempt to radically divorce perceptual experience from bodily agency by appeal to the phenomenology of optic ataxia

It is perhaps unsurprising that the neuroscientists cited above express scepticism over the popularist gloss of the dual stream hypothesis in terms of one information stream for conscious perception (the ventral stream) and another stream for unconscious visually-guided action (the dorsal stream), remarking that it is a something of an oversimplification (Rossetti, Pisella & Vighetto p.178; Pisella et al 2009 p.3042-3). A more thorough review of the empirical evidence against the double dissociation between 'vision-for-perception' and 'vision-for-action' is provided by Shenk & McIntosh (2010) (but cf. Milner & Goodale (2010)). To be absolutely clear, though, I am not dismissing the dual stream hypothesis from the armchair. That would be vulgar, foolish and warrant the full wrath of the neuroscience community. The overwhelming consensus is that some (although probably weakened) version of it is true and as a scientific layman I accept this readily. My point is merely that it is philosophically illegitimate for Clark to invoke empirically contentious details of the dual stream hypothesis in order to carve the sort of radical divide between perception and embodiment needed to undermine anything I have said in this chapter about the role of sensorimotor understanding in shaping ordinary perceptual experience. These concerns are indeed orthogonal to the sensorimotor approach to visual perception.

11. Conclusion

I began by introducing some key aspects of the sensorimotor approach to visual perception, supported by empirical evidence from change and inattentional blindness

experiments. I argued, following Noë and contra some commentators, that change/inattentional blindness does not prove that we are in any way mistaken about the extent of the richness of our perceptual experience. The emerging phenomenological insights concerning the visual field served as the point of departure for a phenomenological analysis of the way in which embodiment shapes our perpetual experience, focussing on the concept of 'sensorimotor understanding' as it is developed in Noë's work and its relevance to what, following Husserl, I called the 'horizonal structure' of conscious perception. I then argued that Noë's account of sensorimotor understanding suffers from certain deficiencies and inconstancies which Merleau-Ponty's early phenomenology equips us to rectify, albeit admittedly at the expense of incurring an inflated conceptual inventory which potentially carries its own distinct set of philosophical problems. Finally, I considered Clark's attempt to empirically discredit the sensorimotor approach by appeal to Milner and Goodale's dual stream hypothesis. These arguments were found to be both empirically suspect and ultimately irrelevant to the phenomenological project of the chapter.

III. Extended Embodiment: Sensory Substitution and the Body Schema

1. Introduction

In the previous chapter, I sought to demonstrate the importance of perceivers' embodiment to the structure of their experience by using Merleau-Ponty's early phenomenology to supplement O'Regan and Noë's sensorimotor approach to vision. An emerging theme of that chapter was that adaptations of the body schema yield novel motor significations and affordances. In this chapter, I further explore the adaptability of the body schema via a phenomenological account of perception via tactile-visual sensory substitution technology. The intended upshot is that by virtue of the body schema's inherent flexibility, it can and sometimes does incorporate items external to the biological body. Doing so will involve delving further into Merleau-Ponty's phenomenology, as well as empirical case studies of sensory substitution technologies along with commentaries from contemporary analytic philosophers. I begin by providing some background information on tactile-visual sensory substitution devices. I then briefly recapitulate Merleau-Ponty's discussion of the body schema before outlining relevant aspects of his accounts of spatiality and the senses, relating these where appropriate to relevant empirical literature. Insights from Merleau-Pontian phenomenology are then applied to perception with tactile-visual sensory substitution devices. Finally, the emerging account is compared and contrasted with those currently available in the literature on sensory substitution.

2. Tactile-visual sensory substitution: some key facts

Sensory substitution devices (hereafter SSDs) are prosthetic technologies designed to enable blind people to perceive the world in ways usually reserved for the sighted (for a scientific overview, see Stiles & Shimojo (in press) and Bach-y-Rita & Kercel (2003)). The focus of this paper is *tactile-visual* sensory substitution (hereafter TVSS). As the name implies, TVSS adapts subjects' surviving sense of touch to achieve a modified form of perception equivalent to vision in many respects. TVSS devices are not the only kind of SSD, though they are the SSDs which have received the most attention from philosophers. There are also audio-visual SSDs which convert camera images into sonic signals, allowing users to experience some of the characteristic phenomenological features of visual perception (see, e.g. Levy-Tzedek et al 2012, Meijer 1992).[29] An adequate phenomenological account of audio-visual SSDs would require a separate treatment far beyond the scope of this chapter (or indeed this thesis), so in what follows I restrict my discussion to TVSS only.

Research into TVSS was spearheaded by the pioneering work of Paul Bach-y-Rita and colleagues in the late nineteen-sixties (Back-y-Rita et al 1969). Their prototype SSD was a relatively crude device which translated images from a video camera into patterns of activation in four hundred small Teflon tipped vibrators (known as 'vibrotactors') fitted to the backrest of an old dental chair which tactually traced the shapes of objects presented to the camera on the user's back. Mappings between camera images and vibratory-tactile stimulation could be monitored as two-dimensional pictures using an oscilloscope. The camera, though bulky and fixed to a tripod, could be manipulated by hand. Six blind subjects (one early blind, five born blind) were trained to use this apparatus. Within a short time they could identify not

---

[29] One such device is 'The vOICe' (capitalised for 'oh I see'), an updated version of Meijer's (1992) original design. Full details about the vOICE including user tutorials and a publication archive can be found at www.seeingwithsound.com (last accessed 14/7/14).

only geometrical forms, but a 'vocabulary' of twenty-five objects including telephones, furniture and toy animals (ibid p.963). The speed with which they identified objects increased with practice and they soon mastered paradigmatically visual concepts including perspective, distance and even lighting (they could detect shadows), which were reflected in their subjective reports (ibid).

Though this primitive SSD was hampered by a lack of portability, it was soon superseded by a range of more convenient devices. One now-iconic device comprised a smaller, head-mounted camera and a sixty-four point array of vibrators which could be attached to the back or abdomen via a belt. Following a short adjustment period (as little a few minutes), users similarly 'learn to make perceptual judgements using visual means of analysis, such as perspective, parallax, looming and zooming, and depth judgements' (Bach-y-Rita 2002 p.500). Amazingly, these newly acquired perceptual abilities are retained when the vibrators are transferred from the back to the front of the user's torso (Bach-y-Rita 2002 p.500). Analogous TVSS devices have since been developed which translate camera images into raised dots or electrical stimulation applied to the tongue (Back-y-Rita et al 1998) or fingertips (Buchs, Maidenbaum & Amedi 2014, Kaczmrek, Tyler & Bach-y-Rita 1997), yielding similar results.

The more accustomed users are to the SSD and the more immersed they are in using it to explore their surroundings, the more recessive their awareness of the tactile-vibratory stimulation becomes (Bach-y-Rita 1972 p.153). Moreover, users react to objects approaching the camera as one would react to objects approaching one's body, as an experimenter discovered accidentally when adjusting the camera's zoom

function while a test subject was still wearing the device. Startled, this subject raised

his arms instinctively to protect his face and lurched backwards away from the

'approaching' object despite being stimulated tactually on his back (ibid. p.98-9). As

well as enjoying many of the advantages of normal vision, blind TVSS users are also

susceptible to visual illusions like the waterfall illusion[30] (ibid. p.79). It is therefore

beyond doubt that TVSS constitutes a genuine form of exteroception

> Our subjects spontaneously report the external localisation of stimuli, in that
>
> sensory information seems to come from in front of the camera, rather than
>
> from the vibrotactors on their back. Thus, after sufficient experience, the use
>
> of the vision substitution system seems to become an extension of the sensory
>
> apparatus. (Bach-y-Rita et al 1969 p.964)

Use of TVSS is a skill – a 'perceptual skill' as Bach-y-Rita often puts it – and as

such requires some adjustment before it begins to feel natural. Bach-y-Rita and his

co-workers realised early on that in order to maximise development of these

perceptual skills, subjects must be able to manipulate the camera for themselves. The

best results occur when subjects manipulate the camera for themselves, with some

experimenters going so far as to claim that active self-movement is necessary for

successful perception (Lenay et al 2003). This was a major drawback of the original

dental chair SSD, where the camera was larger and less mobile. For instance, though

competent with respect to shape, subjects struggled to discriminate sizes (Bach-y-

Rita 1972 p.81). As the sensorimotor approach would predict, the greater capacity for

self-movement granted by subsequent portable SSDs rectified this by allowing for a

---

[30] Motion after effect, or the 'waterfall illusion', occurs when subjects with stationary
eyes are shown a stationary object after being shown a moving stimulus. The
stationary object is then perceived as moving in the opposite direction to the moving
stimulus.

more thorough perceptual exploration of the environment, resulting in more sophisticated visually-guided behaviours. For example, after a short period of adjustment, users of portable SSDs can locate a white object against a black background then walk over to it and pick it up, correctly identify individual characters from a range of letters displayed on a wall (ibid p. 87), and even identify individual faces and facial features. Bach-y-Rita's paraphrase of one user's remarkable report upon exploring a photograph of a member of his laboratory staff using the head-mounted camera testifies to the remarkably vision-like precision of TVSS perception, 'That is Betty; she is wearing her hair down today and does not have her glasses on; her mouth is open, and she is moving her right hand from her left side to the back of her head' (ibid. p.6).

Research into sensory substitution, especially TVSS, has inspired much philosophical debate (for a review, see Kiverstein, Farina & Clark (forthcoming)). In particular, philosophers have been interested in how to characterise the unusual form of perception enjoyed by blind TVSS users. There are three main lines of interpretation of TVSS perception discernible in the literature: the quasi-visual interpretation, the tactile interpretation, and the sui generis interpretation.

Some theorists, including Bach-y-Rita himself, are wont to claim that TVSS perception can be understood as a form of visual perception (Bach-y-Rita 2002; 1972, Hurley & Noë 2003, O'Regan & Noë 2001). This is, after all, the purpose for which SSDs are designed

> If a subject without functioning eyes can perceive detailed information in
> space, correctly localise it subjectively, and respond to it in a manner

comparable to the response of a normally sighted person, I feel justified in applying the term "vision". (Bach-y-Rita 1972 pp.ix-x)

To avoid presenting an overinflated version of the visual interpretation, it must be noted that nobody goes so far as to claim that ordinary vision and TVSS perception are completely phenomenologically indistinguishable; the claim is rather that TVSS perception is sufficiently vision-like to be justifiably described as visual. Hence my use of the term 'quasi-visual'. One obvious difference between ordinary vision and TVSS perception is the absence of colour qualities (at least in the case of users blind since birth, see Ortiz et al (2011)). TVSS perception also typically lacks the affective content that accompanies much visual experience. For instance, pornographic images do not induce sexual arousal, and the perception of loved ones' faces is disappointingly devoid of any affective dimension, though this does not prevent users from taking pleasure in learning to recognise previously unfamiliar objects (Bach-y-Rita 2002 p.509-10).

Despite being a defender of the quasi-visual interpretation, Bach-y-Rita (2002) raises the worry that these facts might count against the possibility of SSDs giving rise to genuinely visual experience. Similarly, Lenay et al (2003) suggest that the lack of emotional value found in TVSS users' experience counts against their experience being genuinely visual. This, I think, is too hasty a conclusion. With respect to colour, blind TVSS users can discriminate coloured surfaces on the basis of light reflectance (Bach-y-Rita 2002 p.502), which is arguably what sighted subjects' functioning eyes allow them to do. Furthermore, we do not deny vision to the colour-blind, nor would we deny vision to a person or creature whose vision was so poor that they could not discriminate colours, or to someone whose vision is entirely monochrome like the painter described by Sacks (1995) who lost his colour vision

due to cerebral achromatopsia. With respect to the absence of emotional content, O'Regan and Noë instructively point out that blind TVSS users' sexual and emotional development occurs in the absence of vision, so it is perhaps to be expected that these experiences will remain non-visual for them (O'Regan and Noë 2001 p.958). And in any case, there exist pathological subjects (such as psychopaths) who lack these emotional attachments altogether: we don't infer from this that they can't see!

Not everyone is impressed by the quasi-visual interpretation. Others have argued that blind TVSS users' perception remains tactile, as the users lack functioning eyes and the source of the perception is tactile stimulation (Prinz 2006, Block 2003, Keeley 2002, Ross 2001). As adaptation to TVSS exploits neural plasticity, further technical disagreement between the two parties arises as to whether the implicated brain processes traffic in visual or tactile information (Hurley & Noë 2003, Block 2003). Against the dichotomy of TVSS-as-visual versus TVSS-as-tactile, others have claimed that TVSS perception is so radically different from both vision and touch that it constitutes an emergent and entirely sui generis form of perception. Deroy & Auvray forthcoming, Auvray & Myin 2009). A related thesis from empirical psychology, which does not immediately contradict any of these, is that TVSS perception consists in an 'artificially acquired' or 'synthetic' form of synaesthesia (Ward and Wright 2012, Proulx & Stoerig 2006). Farina (2013) endorses a combination of the sui generis interpretation and the synaesthesia thesis.

In what follows I undertake a genetic-phenomenological analysis of adaptation to, and perception with, TVSS devices, using further insights from Merleau-Ponty. On a scholastic note, given the influence of Merleau-Ponty's early phenomenology on much contemporary work on embodied cognition, it is perhaps surprising that to date

no such account has been explicitly formulated. I argue that Merleau-Ponty's phenomenology compels us to endorse a hybrid account comprising variations of both the quasi-visual interpretation and the synaesthesia thesis and to reject the tactile and sui generis interpretations, and more generally, that evidence from TVSS further supports my claim in chapter I that the limits of embodiment do not always, and certainly do not by necessity, coincide with the boundaries of the human organism.

3. Adaptation to TVSS via extension of the body schema

Discussions of adaptation to TVSS tend to focus on the plasticity of the implicated neural processing (Block 2003, Hurley & Noë 2001, Bach-y-Rita 1972) and/or the sort of perceptual knowledge gained through the use of an SSD (Prinz 2006, Ross 2001) rather than its phenomenology. Sensorimotor theorists characterise adaptation to an SSD in terms of the acquisition of sensorimotor understanding, and attribute the emerging vision-like perceptual abilities to blind subject's implicit mastery over paradigmatically visual sensorimotor contingencies (O'Regan & Noë 2001 pp.957-8). Again, I am sympathetic to this account, but think that insufficient attention has been paid to the precise ways in which the user's embodiment shapes the adaptation process and the emerging experience. In this section, I argue that blind subjects' adaptation to TVSS consists not *just* in the reorganisation of information processing in their brain and sensitivity to novel sensorimotor contingencies, but in the integration of the SSD with, and subsequent extension of, the body schema.

Recall that the body *schema* is not to be confused with a mental representation of the body - a body *image*. The body schema is an integrated system of motor dispositions

that structures perception and action (Gallagher 2005) and in Merleau-Ponty's words, is a 'law of constitution' (Merleau-Ponty 2012 p.101) in the Husserlian genetic-phenomenological sense of 'constitution' - as coming to appear as such, meaning that it constrains and conditions the ways in which things appear to the perceiver. This pertains both to experience of one's own body and the surrounding world. Indeed, for Merleau-Ponty, the two are inseparable; they are 'two sides of a single act' (ibid. p. 211) which only artificially come apart in 'objective' thought. By this he means that although we can, if we wish, describe in isolation the experience of our own bodies and the experience of the surrounding world, the resulting description would be a piece of abstract philosophical theorising, and not a description of lived experience as it typically unfolds.

This claim motivates Merleau-Ponty to distinguish between 'external' and 'bodily' conceptions of space. External space is space conceived as an unbounded expanse of determinately fixed positions. Bodily space, on the other hand, is a pre-reflective spatial awareness that, for normal subjects at least, is phenomenologically prior to assigning determinate spatial values to objects through an egocentric frame of reference. For example, in order to judge that something is 'about four feet in front of me', I must first have a sense of myself as orientated towards the object through the perspective imposed by my body, without which 'in front of' (and 'me' for that matter) would lack the sense it has in that context. External space and bodily space are therefore complementary notions, as reflective judgements about the former are grounded in a pre-reflective awareness of the latter. Hence Merleau-Ponty's remark that 'Homogeneous [external] space can only express the sense of oriented [bodily] space because it received this sense from oriented space' (ibid. p.104) and his rough-and-ready rebuttal of Kant's transcendental account of space: though in a sense Kant

was correct in holding that an a priori intuition of space is a condition of the possibility of spatial judgements[31], the space Kant has in mind is the universal, objective space of Euclidean geometry and Newtonian physics, rather than the bodily spatiality of unreflective experience

> Insofar as I have a body and insofar as I act in the world through it, space and time are not for me a mere summation of juxtaposed points, and no more are they, for that matter, an infinity of relations synthesised by my consciousness in which my body would be implicated. (ibid. p.140-1)

As illustrated by the case of Ian Waterman in chapter I, the fact that there is no need to 'synthesise' through any deliberative mental act of judgement the individual spatial points at which one's body parts are located into a unified perception of the body owes to the body schema. In the previous chapter we saw, following Merleau-Ponty, that through the honing of skills and acquisition of new habits, the body schema can be 'reworked and renewed' (Merleau-Ponty 2012 p.143) yielding new affordances and motor significations. The body schema is flexible; it is by honing motor habits that we can engage fluidly in skilful behaviour, and the possession of a body schema is a precondition of forming and perfecting such motor habits. Recall also that for Merleau-Ponty, the body schema has an existential significance over and above its relevance to the psychology of perception and action. In fact, it is a key concept in his embodied take on the now-familiar Heideggerian notion of being-in-

---

[31] 'At the basis of their [geometrical objects'] intuition lies a pure intuition (of space and time), which is a *priori*. This is possible because the latter intuition is nothing but the mere form of sensibility, which precedes the actual appearance of the objects, since in fact it makes them possible. Yet this faculty of intuiting *a priori* concerns not the matter of the appearance (that is, the sensation in it, for this constitutes what is empirical), but its form, viz,. space and time.' (Kant 1977 pp.25-6)

the-world. Merleau-Ponty takes up the Heideggerian notion of being-in-the-world and gives it an embodied twist: the body is the 'vehicle of being-in-the-world' (Merleau-Ponty 2012 p. 84), because it is through the body schema that we have a practical grasp on our surroundings prior to deliberation or reflection. On Merleau-Ponty's account, even instinctual postural adjustments like the hand movements implicated in grasping, or avoidance behaviour induced by pain, are more than blind stimulus-response reflexes because they have a 'sense'; they are the appropriate bodily adjustments given the implicitly understood normative demands of one's situation, 'Reflex, insofar as it opens itself to the sense of a situation, and perception, insofar as it is an intention of our total being, are modalities of a *pre-objective perspective* that we call "being in the world"' (ibid. 2012 p.81, italics in original). Hence, to inhabit the practical sphere that Heidegger emphasises as primary for Dasein, one must be anchored into the world through a body schema.

Recall also Bach-y-Rita's remark that SSDs are 'extensions of the user's sensory apparatus'. With this in mind, Merleau-Ponty's account of the body schema allows us to make phenomenological sense of blind subjects' adaptation to TVSS SSDs: TVSS *extends* the body schema beyond the boundaries of the biological body and functions as part of the system of embodied capacities perceivers bring to bear on their environment through perceptual exploration. As Bach-y-Rita observes, it is essential to successful use of an SSD that 'camera movement must be under the control of one of the subject's motor systems (hand, head movement, or any other)' (Bach-y-Rita 2002 p.497) in order to compensate for the absence of the exploratory eye movements, and indeed other bodily movements such as head turnings, upon which ordinary vision is heavily dependent (Bach-y-Rita 1972 p.99, cf. Yarbus 1967). Recall also that, as with awareness of one's body generally, awareness of the tactile

stimulation is highly recessive. By developing the required motor habits, the SSD ceases to be given as an external object for the user, and becomes what Heidegger (1962) calls 'transparent equipment', or what Rowlands (2010) has more recently dubbed a 'vehicle of disclosure': it becomes for them something *through which* they perceive the world, rather than a mere object which that they perceive as out there in the world.

As with Head and Holmes' example of the blind man's cane noted in chapter I, the SSD becomes integrated into the user's bodily space, and so long as they are immersed in the exploration of their environment the vibrators on their body, like the handle of the cane, are not objects of perception, but part of the pre-reflective system of motor capacities though which perceivers are aware of the surrounding world. However, unlike the blind man's cane, which extends his field of tactile sensitivity, perception with an SSD allows the user to experience objects with which their sensory apparatus is not in direct physical contact. The objects of TVSS perception, - faces, walls, pictures, etc - are no more experienced at the source of the vibratory simulation than the objects of ordinary vision are experienced as on the retina, but at their actual external locations. It is in this crucial sense that blind TVSS users' newfound sensory fields have more in common with vision than they do with touch because, as Martin (1993) notes, while one's field of tactile sensitivity coincides with the boundaries of one's body (phenomenal space), the visual field appears as distal and unbounded (external space). Hence, adaptation to TVSS allows blind subjects gain the sort of *perspective* on external space that is usually provided by vision. This notion of perspective harkens back to Merleau-Ponty's description of the visual field:

> To say that I have a visual field means that I have an access and an opening to
> a system of visible beings through my position, and that they are available to

my gaze in virtue of a kind of primordial contact and by a gift of nature,

without any effort required on my part. (Merleau-Ponty 2012p.224)

Adaptation to TVSS doesn't just enable blind users to identify an inventory of items learned by rote, but to grasp motor significations and thereby respond intelligently and appropriately to their situation via the exercise of sensorimotor understanding. It is by acquiring a quasi-visual field that blind TVSS users can orient themselves in relation to distal objects and anticipate bodily engagements with them. The user who lurched backwards had neither the time to make a judgement about the looming object, nor a prior history of 'zooming' stimuli from which to develop a conditioned reflex[32], but having assumed the perspective of the camera, the sudden stimulation was instantaneously *understood.* This practical form of understanding is an achievement not of a reflecting subject, but of the pre-reflective operative intentionality of the body schema.

4. Aside: A quick note on Kiverstein and Farina on SSDs and extended consciousness

At this juncture, it is worth noting how my Merleau-Pontian account of adaptation to TVSS (which at this point in the chapter is, admittedly, far from complete) differs from a prima facie similar one due to Kiverstein and Farina (forthcoming). They argue that SSDs form part of an extended (i.e. not entirely neural) physical basis for their uses' conscious experience. With this much I agree. Indeed, a defence of the 'extended mind thesis' is the subject of my next chapter, and a defence of what I shall

---

[32] In the terminology of *Phenomenology of Perception,* these alternatives would be labelled respectively as 'intellectualist' and 'empiricist'.

call the 'extended conscious mind' thesis is the theme of chapter V. I am also in agreement with their claim that this extension of the physical states and processes which realise the SSD user's perceptual experience is made possible by an adaptation of their body schema. However, in contrast to the Merleau-Pontian picture I have presented, Kiverstein and Farina follow neuroscientists in characterising the body schema as ' an action-oriented [neural] representation of the body that gives the agent knowledge of what he can do with his body' (Kiverstein and Farina 2012 p.21). Hence, despite their aim of defending externalism about consciousness, Kiverstain and Farina apparently have an internalistic take on the body itself. If adaptation to TVSS consists merely in the forming of a novel kind of internal representation, then it is unclear why anything outside of the SSD user's brain should count as part of the physical substrate of their experience[33]. If the body is merely a causal pathway through which the requisite internal representations come to be formed, then Kiverstein and Farina need to show why the SSD itself should amount to more than an extra component in the causal chain or fall foul of what Adams and Aizawa (2001; 2007) call the 'coupling-constitution fallacy' (I discuss this fallacy in more depth over the next two chapters). By construing the process of adaptation to TVSS as an adaptation of the sensorimotor mechanisms which comprise the body schema proper, and not reducing these to internal representations thereof, the Merleau-Pontian account is actually far better suited to the cause of externalism about consciousness.

5. The unity of the senses, synaesthesia, and bodily synthesis,

---

[33] Nb. Again, I am not denying that the forming of novel action-oriented representations is involved in adaptation to TVSS, only that this tells the whole story.

Thus far in following Merleau-Ponty I have claimed that TVSS perception is made

possible by the incorporation of the SSD into the body schema, and that users

thereby gain a perspective on objects in external space via the acquisition of new

sensorimotor understanding. In this section, I turn to part two of *Phenomenology of

Perception* to shed further light on the experience of TVSS perception.

In the 'Sensing' chapter of *Phenomenology of Perception,* Merleau-Ponty argues the

intentionality of the senses is not reducible to that of isolated sensory systems. This is

because sensory experience is governed by an underlying operative intentionality,

whereas attending to the contribution of individual sensory modalities involves a

form of object intentionality. The justification for this claim stems from the simple

phenomenological observation that perceptual experience does not typically present

us with pre-distinguished visual qualities, tactile qualities, etc, but unified inter-

sensory wholes whose relation to a particular sensory modality is often ambiguous.

The question of how to distinguish sensory modalities is the subject of longstanding

debate in the philosophy of mind (for a logical geography see Sorabji (1971), Grice

(1962), Leon (1988), Ross (2001), Keeley (2002), and Macpherson (2010))[34], and

although it is not Merleau-Ponty's aim to discredit this enterprise, he is critical of

---

[34] Leon (1988) argues that the senses are distinguishable by their unique, intrinsic
phenomenal characters or 'qualia'. The classic formulation of this view is generally
attributed to Grice (1962), though it is actually one of several criteria Grice discusses
in his treatment of the senses, albeit the one on which he places the most emphasis.
Against the qualia view, Ross (2001) argues that the senses are to be distinguished by
reference to the particular properties of objects they detect. This 'proper objects' view
has its origins in Aristotle's *De Anima,* as Sorabji (1971) makes clear. Keeley's
(2002) account of the senses, which demarcates sensory systems in terms of their
proper biological functions, can be read as an empirically-informed naturalistic
version of the proper objects view. Macpherson (2010) questions the extent to which
these approaches are mutually exclusive, and develops a 'multidimensional space'
account whereby different approaches to individuating sensory modalities each
contribute to a wider project of delimiting possible (e.g. animal, robotic, alien)
senses.

traditional philosophy and psychology's attempts to 'pulverise' (Merleau-Ponty 2012 p.236) experience into sense-specific qualities or representational contents. This is because performing the mental experiment of segregating the senses requires the adoption of a somewhat abstract and psychologically unnatural 'analytic attitude' (ibid p.247) towards ordinary experience in response to a given question which presupposes their distinctiveness, whereas ordinary experience is inherently *multisensory.* Merleau-Ponty therefore sides with a growing number of analytic philosophers who express scepticism over the extent to which sensory experiences are always or even ever modality specific (see e.g. the essays collected in Stokes, Matthen & Biggs forthcoming). Philosophers therefore risk committing the 'experience error' of producing a distorted description of perception in terms of what we come to believe subsequently about its objects through commonsense reflection or scientific investigation, at the expense of accurately describing the phenomena under investigation (Merleau-Ponty 2012 p.5). By means of illustration, consider the way by which Husserl arrives at his notion of 'pure sensation'. Husserl writes that sensory experience consists in the 'apprehension' of sensations through which the 'schemata' of objects are constituted. In other words, perception presupposes sensory inputs which are organised into unified appearances of objects, 'Perceptual apprehension presupposes sensations-contents, which play their necessary role for the constitution of the schemata and, so, for the constitution of the appearances of real things themselves' (Husserl 1989 p.61). However, in contradistinction to the broadly Kantian picture implied by this claim, Husserl does seem to take it that these raw sensations can feature directly in the content of perceptual experience, as illustrated in his discussion of the 'presupposed sensation-contents' in the auditory perception of a tone

Finally, the spatial apprehension could be suspended, and then it becomes a mere "sense-datum" instead of a spatially sounding tone […] It is to be understood that such a tonal datum could be constituted without there being carried out any spatial apprehension at all, an apprehension which, in our example, was only abstractly put aside or, to correct this misleading expression, was suspended but is still, in the changed mode, a lived experience, the lived experience, to be precise, which pre-gives the spatial tone. We should mention, however, that this is not a necessary pre-givenness. A tone would be thinkable which dispensed with every spatial apprehension. Here, with the pure datum of sensation, we encounter a pregivennes which yet precedes the constitution of the object as object." (ibid. p.24-5)

Here Husserl flip-flops between two claims. One the one hand, he introduces the notion of pure sensation as the product of a scrutinising attentive attitude to one's current sensory experience. Husserl does, after all, start by saying that the spatiality of the perception *could* be subtracted away from the experience. By the end of the passage, though, he appears to slip into a second claim to the effect that by hypothetically abstracting away the spatial organisation of the perception, we uncover a pure sensory datum which *was* previously present, albeit unnoticed, in the experience, by inferring that the phenomenologist's inspection of her sensory experience directs her towards a 'pre-given' 'pure datum of sensation' that preceded her perception of the object. In other words, Husserl here endorses what might be called a pre-reflective sense datum theory of sensation.[35] Husserl's account of pure sensation therefore presents the phenomenologist with a dilemma. Either the notion

---

[35] This position, phenomenalist in spirit, coheres with the recurring notion of 'hyletic data' in Husserl's account of the lived body in *Ideas II* (Gallagher 1986b, cf. Q. Smith 1977).

of a pure sensation is an abstraction from our ordinary perceptual encounters with the world, in which case it is an inappropriate posit for a phenomenological description of ordinary sensory experience, or object perception is consequent upon unnoticed pure sensations akin to atomistic sense data. Merleau-Ponty takes the first option by, unlike Husserl by the end of this passage, remaining true to the 'appearances of real things themselves'

> Far from being coextensive with perception, the sensible quality is the peculiar product of the attitude of curiosity or observation...[T]he quality does not figure in the natural exchange between my vision and the world. It is the response to a certain question posed by my gaze and the result of a second-order or critical act of vision that attempts to know itself in its particularity; it is the result of an "attention to the purely visual" that I employ when I am worried about being tricked or when I wish to commence a scientific study of vision. (Merleau-Ponty 2012. p 235)

Against the abstract notions of 'purely visual' and 'purely tactile' contents, Merleau-Ponty stresses the extent to which sensory modalities are integrated, yielding a holistic and multisensory perceptual field or, as he puts it, the way in which they 'communicate between themselves' (ibid. p.238). An especially relevant example, which appears twice in *Phenomenology of Perception,* is the 'communication' between touch and vision in the perception of a furry carpet. Merleau-Ponty remarks, 'this red would literally not be the same if it were not the "wooly [sic] red" of the carpet' (ibid. p.5), and later, 'the blue of a rug would not be the same blue if it were not a wooly [sic] blue' (ibid. p.326). That he invokes the same example using different colours is not insignificant, for Merleau-Ponty's point is not that perceivers can associate the carpet's colour with the way they expect it to feel, but that

irrespective of its particular colour or what they happen to know about the purpose of carpets, they *literally see its texture*; it is *seen* as *touchable*.

For Merleau-Ponty, this sort of inter-sensory communication is phenomenologically ubiquitous. Consider, for example, how walls *look* solid, or how gravel paths *look* rough (see O'Callaghan (forthcoming) for additional examples). The horizonal structure of these perceptions include non-visual co-present contents. Similarly, certain colours are experienced as 'warm' or 'cool', and evidence strongly suggests that these sorts of correlations between colours and affective-sensory states hold across cultures (Gao et al 2007, Ou et al 2004). In instances like these, we are only tempted to factor the experience into uniquely visual and tactile contents when we cease to describe experience qua experience, and commit the experience error. Phenomenologically, what is sensed is not a pure visual given, but a multisensory whole which is simultaneously visible and touchable, and embedded in an implicitly understood context of potential bodily interactions which Merleau-Ponty calls a 'motor physiognomy' (Merleau-Ponty 2012 p. 217). Although we can selectively attend to the 'distinct but indiscernible' contributions (ibid. p.239) of individual sensory systems to the overarching perceptual experience, this cannot destroy perception's multisensory structure by reducing experience to an isolated sense-specific quale or sense-datum[36].

---

[36] Note that it makes no difference to Merleau-Ponty's phenomenological claim whether the metaphysics of these hypothetically isolated sensory contents are conceived in terms of sense-data-like mental intermediaries, a la the qualia view, or the properties of objects they represent, a la the proper objects view. Tye (2002), for example, argues rather convincingly that the former are reducible to the latter. All that matters here is that these do not show up as pre-distinguished in ordinary perceptual experience.

Of course, appearances can be deceiving; it could transpire that the 'solid' wall is actually a hologram, that the 'rough' gravel is made of carefully painted pieces of smooth polystyrene, and so on. But this is an epistemological side issue that does not undermine this phenomenological claim. And the senses can conflict, as happens with the McGurk effect noted in chapter I, but while conflict between senses may cause confusion, it is only against an assumed background of general harmony that this conflict can appear surprising or troublesome, and the fact that this harmony can be disturbed does not mean that it does not usually hold. In any case, epistemologically undesirable communication between sensory modalities is still communication between sensory modalities. This emphasis on the multisensory nature of the perceptual field allows us to make sense of Merleau-Ponty's otherwise obscure claims about the phenomenon of synaesthesia, which in turn has ramifications for a phenomenology of TVSS perception.

In the psychological literature, synaesthesia is generally regarded as a perceptual abnormality (though not a disability) where a phenomenal feature typically associated with one sensory modality is experienced idiosyncratically in another modality. Until relatively recently, reports of synaesthesia were largely written off as fabrications, loose metaphorical talk or effects of hallucinogenic drug abuse (LCD and mescaline can induce similar experiences). However, scientific consensus has now shifted towards universal acceptance of the phenomenon as genuine, caused by cross-wiring between cortical regions (Ramachandran & Hubbard 2001a). It is likely to have a genetic basis (Ward & Simmer 2005), is estimated to occur in around four to five per cent of the population (Simmer et al 2006), and despite the once widespread belief that it is 'overwhelmingly a female condition' (Baron-Cohen 1996 p.1), ongoing research suggests it effects persons of both sexes equally (Simmer et

al, ibid). The most studied form of synaesthesia is grapheme-colour synaesthesia, where subjects see letters or numerals as coloured (Ramachandran & Hubbard ibid, Galton 1881). Other forms of synaesthesia include auditory-visual synaesthesia, where sounds automatically induce reliable visual experiences (Goller, Otten & Ward 2008), taste-touch synaesthesia, where contact with certain textures automatically induces particular tastes (Ramachandran & Hubbard 2003), and tactile-visual synaesthesia, where tactile stimulations automatically induce reliable visual experiences (Simmer & Ludwig 2012). For present purposes, though, we need only to distinguish broadly and heuristically between what we may call *intra-modal* and *inter-modal* forms of synaesthesia. Grapheme-colour synaesthesia is intra-modal, because colours and graphemes are both paradigmatically visual contents. More relevant to the current discussion is inter-modal synaesthesia, where the sensory contents typically associated with one modality are experienced via a different modality, as occurs in audio-visual, taste-touch and, most importantly in the present context, tactile-visual, synaesthesia. It is this kind of inter-sensory communication between vision and touch that Merleau-Ponty has in mind with his example of the furry carpet.

For Merleau-Ponty, inter-modal synaesthesia is not radically different from ordinary perception, but a modified form of experience where certain features of ordinary perception are disproportionately amplified or recalibrated. What is abnormal about the cases of inter-modal synaesthesia documented in the psychological literature is therefore not the fact that they are synaesthetic, but their particular idiosyncrasies and the strength with which they are experienced. That is to say, there is no synaesthesia *ex nihilo*, because the senses already interact to yield unified, multisensory perceptual experience

Synaesthetic perception is the rule and, if we do not notice it, it is because scientific knowledge displaces experience and we have unlearned seeing, hearing and sensing in general in order to deduce what we ought to see, hear, or sense from our bodily organisation and from the world conceived by the physicist. (Merleau-Ponty 2012 p.238)

At various points in *Phenomenology of Perception* and elsewhere (e.g. Merleau-Ponty 1964d) Merleau-Ponty appeals to empirical studies of infants to support claims about lived experience prior to its distortion in 'objective' theorising. It is fitting, therefore, that his claim that perception is originally synaesthetic is supported by recent work in developmental psychology. Earlier psychologists, notably Piaget (1952), believed that infants' sensory systems initially function independently, and are subsequently integrated through learning. Against this view, there is neurological evidence to suggest that infants' perception actually involves a substantial amount of cross-modal activation between sensory systems. Though neural connections between sensory areas of the brain are 'pruned' in normal development, some are retained into adulthood (Spector & Maurer 2009, Maurer & Mondloch 2005, Maurer & Muarer 1988). The presence of inter-modal synaesthesia in infants is thought to play an important role in their perceptual and linguistic development (Maurer, Pathman & Mondloch 2006, Baron-Cohen 1996, Köhler 1947), as well as the cultural development of language and its evolutionary origins (Ramachandran & Hubbard 2001b, Day 1996). Despite their differences in aims and methodology, the moral of these empirical findings is distinctly Merleau-Pontian: perception is inherently multisensory and ambiguous with respect to specific modalities, even in normal adults, but we typically fail to appreciate this in our analyses because of a

tendency to substitute an idealised rational reconstruction for an accurate description of the phenomena.

It is Merleau-Ponty's contention that the integration of sensory modalities is also accomplished by the body schema, through which they are unified into a singular act of an embodied subject. He explains this using an analogy with double vision. Pushing one's eyeball causes a retinal disparity resulting in a double image. Upon release of the eye, the two images merge. The merging of images is not achieved through a mental act of judgement; it simply unfolds 'before our eyes' (as it were). In this sense, the merging of images is a mechanical operation of the eyes and wider visual system. However, considered phenomenologically, the merging of images is more than mechanical because it has an existential significance for the subject: the double image is experienced as an *obstacle* to ordinary perception to be transcended. The merging of images is lived through as a progression from a confused perceptual state which violates a certain norm of appearance, towards a *corrected* or *completed* perception.[37] The double image poses a problem for the perceiver, but this problem is resolved pre-personally (as Merleau-Ponty would put it) by a body structured in such a way as to comport itself towards the completion of perception

> And herein lies the distinction between perceptual synthesis and intellectual synthesis. When I pass from double vision to normal vision, I am not merely conscious of seeing the same object with my two eyes, I am also conscious of progressing towards the object itself and finally of having its carnal presence.
> (Merleau-Ponty 2012 p.242, italics removed)

---

[37] Husserl too makes this observation, and it is likely that here Merleau-Ponty is channelling his remark that double images 'mean things only in contradiction to all normal motivations' (Husserl 1989 p.66, italics removed).

On Merleau-Ponty's account, the same goes, mutatis mutandis, for sensory modalities, which are unified through neither 'transcendental' judgement (a la his 'intellectualist') nor 'empirical' association (a la his 'empiricist'), but through a 'bodily synthesis' (ibid. p.154) whereby the 'distinct yet indiscernible' senses are intentionally directed towards one and the same environment

> The vision of sounds or the hearing of colours comes about in the same way as the unity of the gaze through the two eyes, insofar as my body is not a sum of   juxtaposed organs, but a synergetic system of which all of the functions are taken up and tied together in the general movement of being in the world, and insofar as it is the congealed figure of existence. (ibid. p.243)

This notion of the sensing body as a 'synergetic system' sheds further light on perception via TVSS. Sensory experience is 'synergetic' in the literal sense of being more than the sum of its parts. It is not reducible to the qualities or representational contents of its composite subsystems, because the objects of perception are not isolated sensory givens but multi-sensible objects to which the senses, organised into a unified bodily system, are uniformly intentionally directed. It follows that TVSS perception is synaesthetic, but only insofar as *all* perception is, in a limited and qualified sense, synaesthetic, because sensing is a singular act of a bodily system comprising multiple integrated sensory subsystems and not a motley of juxtaposed acts of seeing, touching, etc. Hence, like the experiential differences between diagnosed synaesthetes and normal perceivers, the experiential differences between the sighted and blind TVSS users consist in the *organisation* of their perceptual field, 'The blind person's world and the world of the normal person differ not merely in the quantity of matter available to them, but moreover in the *structure* of the whole' (ibid. p.233, italics in original). A Merleau-Pontian phenomenological account of

TVSS therefore agrees with Farina (2013), Ward and Wright (2012), and Proulx and Stoerig (2006) that TVSS perception is synaesthetic, but denies that this synaesthesia is 'acquired' or 'artificial'. A degree of inter-modal synaesthesia is inherent in perceptual experience, and is therefore not 'acquired' in the sense in which knowledge or a medical condition is 'acquired', and is not 'artificial' or 'synthetic', but rather quite natural. Adaptation to TVSS does not consist in the acquisition of inter-modal synaesthesia, rather, the prior presence of low-level inter-modal synaesthesia in the sensory field, due to unpruned (sic) neural connections and the bodily system in which they are embedded, makes adaptation to TVSS possible. This seems particularly plausible given the presence of unlearned tactile-visual synaesthesia in blind subjects who do not use TVSS devices (Steven & Blakemore 2004). TVSS merely exploits the extant but more limited multisensory perceptual field enjoyed by blind subjects by endowing it with a quasi-visual perspective.

Once it is acknowledged that in ordinary experience objects are *seen as touchable*, we can understand the role played by blind TVSS users' surviving sense of touch without contradicting their firsthand reports by assimilating their experience to a form of extended tactile sensitivity. Recall that users experience the objects of TVSS perception as externally located, relative to an embodied perspective, and that their awareness of the vibratory stimulation form the SSD is highly recessive. To have an object in one's visual field is, at least in part, to recognise its touchability (sic) from a distance. There is therefore a tactual component to their experience, but it is the same tactual component present in ordinary vision. Once integrated, the SSD functions as part of the bodily synthesis by which sensory stimulations are unified into a holistic perceptual field, and what is restored to the blind through TVSS is not a set of uniquely visual qualities or concepts, but the capacity to grasp the motor

signifations that externally located objects - surfaces, faces, familiar items, etc -
have for an embodied, situated subject, that is, a subject with a perspective.

6. The Merleau-Pontian account vs. the alternatives

To recapitulate, I have argued that TVSS perception should be understood in terms
of an extended singular bodily act of sensing made possible by the integration of the
SSD with the user's body schema, and that the perceiver thereby acquires a quasi-
visual perspective allowing them to sense the touchability of objects in external
space. The Merleau-Pontian account therefore endorses elements of both the quasi-
visual interpretation and the synaesthesia thesis. The truth of the quasi-visual
interpretation consists in the fact that the enhanced sensory field enjoyed by blind
TVSS users has more in common with a visual field than an area of tactile
sensitivity, because its objects are given as located in external space relative to the
user's own bodily space, which now incorporates the SSD. The truth of the
synaesthesia thesis is that TVSS perception is indeed synaesthetic, but this is not to
be explained as the acquisition of synaesthesia ex nihilo caused by interaction with
the SSD, but by endowing the subject's inter-modally synaesthetic sensory field with
a quasi-visual perspective. In this final section, I situate the Merleau-Pontian
account in relation to extant interpretations of TVSS perception. I begin by arguing
that the tactile and sui generis interpretations should be rejected. I then highlight
some ways in which the Merleau-Pontian account avoids certain problems incurred
by versions of the quasi-visual interpretation as articulated in the existing literature
on sensory substitution.

From a phenomenological point of view, the tactile interpretation is a non-starter. The fact that it is starkly at odds with how skilled TVSS users typically describe their experience is reason enough to treat it with a degree of scepticism, but several philosophers have insisted nevertheless that TVSS perception is indeed tactile. In order to contradict blind TVSS users' first hand reports, defenders of the tactile interpretation tend to resort to a sub-personal level of explanation which bypasses phenomenological description entirely in favour of attending to perception's underlying physical causes. For instance, Keeley's (2002) insistence that TVSS perception is tactile follows from his view that sensory modalities are to be individuated by reference to proper functions of dedicated sensory systems as determined by natural selection, while Block (2003) appeals to neuroscientific data on the processing of spatial information in the somatosensory cortex to support the claim that the information processed by the brain while using an SSDs remains tactile. But not only does this presuppose a dichotomy of purely visual versus purely tactile contents we have grounds to treat with suspicion, it is a latent instance of the experience error, and therefore inadequate for the task of understanding perceptual experience.

Another tactic for establishing the tactile interpretation, which is less deferential to facts extrinsic to phenomenology, is the appeal to a two-step reasoning process allegedly employed by blind TVSS users. Ross (2001) makes such a move. He distinguishes between 'direct' and 'indirect' perception. Not to be confused with the traditional philosophical antitheses of direct (or 'naive') and indirect realism about the objects of perception, in Ross' terminology 'direct' perception is non-inferential and 'indirect' inferential. However, he adds the following qualification, 'indirect perception can be non-inferential, if by 'non-inferential' we mean that there is no

reflective inference; that the couch looks soft can become non-inferential in this sense' (ibid p.501, my italics), and interprets TVSS perception as one such case of the need for reflective inference being alleviated by 'strength of association' while continuing to count as indirect (ibid.). This empiricist sentiment is echoed by Prinz: 'My best guess is that prosthetic vision devices simple [sic] allow subjects to make automatic inferences about where objects are located in space as a result of tactile information' (Prinz 2006 p.5). The claim is therefore that TVSS perception involves *unconscious inferences* to visible properties based on information gained through touch. The alleged role of unconscious inferences in perception, due originally to Helmholtz (e.g.1962), is the subject of some controversy (Hatfield 2002 cf. Fodor & Pylyshyn 1981). From a phenomenological point of view, however, the claim that TVSS perception involves inferences based on tactile input is straightforwardly false. By construing the required inferences as unconscious, Ross and Prinz tacitly acknowledge that the *experience* cannot be adequately described in inferential terms. Although it is true that TVSS users make inferences in the learning stages, often taking several seconds or longer to work out what they're perceiving (Bach-y-Rita 1972 p.6), the claim that developed TVSS perception consists in rapid inferences is only plausible given the assumption that adaptation to TVSS consists solely in learning and honing such inferences. This over-intellectualises the process, which, as we have seen, is more about honing motor skills and grasping motor significations than exercising judgements. To perceive with an SSD is to exercise a practical skill, and practical skills are not gained simply by sharpening one's reasoning, but through bodily habituation; I might be able to recognise the shapes of E, A and D chords on a guitar, but I won't be able to properly play a twelve-bar rhythm track until my fingers are disposed to shift seamlessly between them. Furthermore, this gloss of

TVSS perception in terms of a fictitious two-step reasoning process misses the vital phenomenological point that the perceiving subject is not a disembodied reflecting ego, but an embodied subject at home in the world. Trivially, TVSS perception involves tactile stimulation, but this does not make it a form of tactile perception.

Merleau-Ponty's phenomenology also compels us to reject the sui generis interpretation of perception with a tactile-visual SSD. Within the Merleau-Pontian framework outlined above, the claim that TVSS perception constitutes an entirely sui generis form of perception seems unnecessarily radical. While all sides can agree with the relatively innocuous claim that SSDs 'can become integrated to such a degree that one can speak of a novel form of perceptual/sensorimotor interaction' (Auvray & Myin 2009 p. 1053), we should stop short of the far stronger claim that through TVSS 'perception can not only transcend its sensory origins, but it can move beyond the confines of the traditional senses' (ibid. p.1054). This only seems compelling if 'traditional' sensing is conceived in terms of isolated visual acts, tactile acts, etc, rather than as singular act of a synergetic bodily system. If we follow Merleau-Ponty in holding that even ordinary sensory experience cannot be adequately described in these terms, then the fact that TVSS perception cannot be adequately captured in these terms either does not license the inference that it constitutes an entirely novel form of perception distinct from the natural senses. And while it is beyond doubt that ordinary perception and TVSS perception differ experientially, inferring from this that 'these emerging capacities brought by SSDs do not and cannot figure on the same list as the natural or existing senses' (Deroy & Auvray forthcoming p.17) is one step too far. There is enough overlap between the two forms of perception to render the sui generis interpretation unnecessary; their difference is one of structure, not of kind. Furthermore, the sui generis interpretation

risks rendering TVSS perception mysterious in a deeply undesirable way. As Lenay et al (2003) stress, we should not neglect the important fact that SSDs increase the extent to which the blind and the sighted partake in a common experience of a shared world. By emphasising the bodily being-in-the-world common to both, the more conservative Merleau-Pontian account allows us to recognise their shared intentional structure thereby avoiding the potential epistemological and ethical quandaries incurred by rendering blind TVSS users' experience radically alien, whilst continuing to acknowledge their genuine phenomenal differences.

Finally, we can see how the Merleau-Pontian account improves on previous attempts to defend the quasi-visual interpretation of TVSS perception. On a negative note, it is instructive with respect to how *not* to argue for the quasi-visual interpretation. Any neuroscientific argument, such as Hurley and Noë's (2003) which appeals solely to neural plasticity[38], is off the table for the same reason that Block's argument for the tactile interpretation is - facts about the implicated neural activity and the correct characterisation of the information processing therein, while scientifically interesting and valuable, cannot serve as a reliable guide to a phenomenological analysis on pain of committing the experience error. Such approaches presuppose, rather than provide, an adequate phenomenological description from which their characterisation of the neural processing, as visual or tactile, derives.

Another possible line of argument, with which Bach-y-Rita (2002), Hurley and Noë (2003) and Noë and O'Regan (2001) all flirt, is also clearly ruled out by the Merleau-

---

[38] Hurley and Noë distinguish between cortical dominance and cortical deference, and argue that TVSS perception is a form of the latter. 'In cases of *cortical dominance*, cortical activation from a new peripheral input source gives rise to experience with a qualitative character normally or previously associated with cortical activity in that area...In cases of *cortical deference*, in contrast, cortical activity in a given area appears to take its qualitative expression from the character of the nonstandard or new input source. (Hurley & Noë 2003 p.133)

Pontian account: appeal to uniquely visual qualities or objects. It is no coincidence that arguments offered in favour of both the tactile and quasi-visual interpretations of TVSS perception tend to derive from a (sometimes explicit, sometimes implicit) logically prior account of sensory modality individuation. Proponents of the qualia view of sensory modality individuation (to which Bach-y-Rita (2002) leans), should they wish to argue for the visual interpretation, would need to identify a unique mind-dependent sensory quality present in ordinary vision, absent from the other senses, that is restored by TVSS. Proponents of the proper objects view of modality individuation are saddled with the analogous task of singling out a particular property or class of properties of objects to which vision alone gives access that TVSS devices allow the user to experience. By following Merleau-Ponty in rejecting as phenomenologically wrongheaded the idea of introspectively isolatable sensory qualia and the dichotomy between purely visual and purely tactile contents, the defender of the quasi-visual interpretation is disburdened both from identifying such a quale and from answering the question of which if any properties are exclusively visual. On a positive note, the Merleau-Pontian framework does allow us to make sense of the precise way in which TVSS perception is quasi-visual without the problematic appeal to sense-specific qualities or contents, by emphasising the uniquely visual *perspective* with which users are endowed by virtue of an extension of the body schema. It is *this* phenomenological feature of TVSS perception that justifies the characterisation of the experience as quasi-visual rather than tactile or sui generis.

7. Conclusion

I began by describing some classic examples of tactile-visual SSDs and detailing the ways in which these allow blind subjects to perceive the world in ways approximating to ordinary vision. I then appealed to further insights from Merleau-Ponty's *Phenomenology of Perception* to develop a novel and plausible account of TVSS perception which avoids some serious problems incurred by the available alternatives, to explore one way in which the body schema can adapt to incorporate objects beyond the boundaries of the organism. On the account of TVSS I have developed, the SSD functions as an external part of the same 'bodily synthesis' by which stimulations of sensory subsystems are integrated into a unified and multisensory perceptual field. Adaptation to TVSS consists in the integration of the SSD with the subject's flexible body schema via the honing of sensorimotor understanding. The acquisition of a quasi-visual perspective, and the consequent ability to engage in visually-guided behaviours, is understood in terms of a modification of bodily space which enables the subject to orient themselves in relation to objects in external space. It is by virtue of its inherent flexibility that the body schema does not by necessity coincide with the boundaries of the human organism. I also argued, contra certain detractors, that the emergent experience, though dependent on tactile stimulation and qualitatively different from the vision of the sighted, can nevertheless still be understood as a limited form of vision.

IV. Beyond the Body: Extending the Mind

1. Introduction

In the previous chapter, I presented the case for a form of 'extended' embodiment via sensory substitution. This chapter branches even further outwards to explore the significance of items located beyond the boundaries of the body entirely to defend the clam that many instances of cognition - reasoning, action-planning, remembering - do not take place solely inside the brains of intelligent agents, but are constituted by physical process which loop brain, body and world. In other words: the mind is not (always) in the head.

The belief that the mind is located inside the head is so longstanding and widely held that for many working in philosophy and the cognitive sciences it is practically axiomatic. Indeed, the central question for philosophy of mind in the analytic tradition, the vast majority of which is predicated on some version of physicalism, is that of how, exactly, mental states can be caused by, supervene on, or be identified with, states of the brain.[39] Even Descartes, the historical poster child for substance dualism, did not doubt that immaterial mental substance is housed inside the skull, interacting with the organism via the pineal gland of the brainstem (Descartes 1985 pp.99-108). Over the last two decades, however, this default assumption has been the target of an ever-expanding wealth of literature advocating and defending a position variably known as 'active externalism' (Clark and Chalmers 1998), 'the hypothesis of extended cognition' (Clark 2008), 'vehicle externalism' (Hurley 2010) or more

---

[39] See, for example, the way in which the debates over the different brains of physicalism, functionalism, dualism, eliminativism, etc, are framed in Kim (2006) and the papers collected in Rosenthal (1991).

generally and in the idiom I shall adhere to, *the extended mind thesis*. According to the extended mind thesis, thinking is not always something we do entirely in our heads. Quite often, thinking quite literally takes place in the 'external' world. To be more precise, extended mind theorists claim that a good deal of cognitive processes - the physical processes which constitute mental phenomena like reasoning and the intelligent guidance of action - comprise not *just* states of the brain and central nervous system, but *also* parts of the environment in which the organism that houses the brain is embedded. The extended mind thesis is attracting a growing number of adherents (e.g. Rowlands 2012, Clark 2008; 2003; 1997, Wheeler 2007, Menary 2007a, Wilson 2004, Clark and Chalmers 1998, and the essays collected in Menary 2010a, cf. Rupert 2009, Adams and Aizawa 2007; 2001), and its converts have explored the possibility of its application to areas as diverse as social cognition (Krueger 2012, Krueger & Overgaard 2012), the psychology of memory (Sutton et al 2010), personal identity (Olsen 2011, Rudder-Baker 2009), social ontology (Gallagher 2012, Gallagher & Crisafi 2009), musical performance (Krueger 2014, Cochrane 2008) and the study of language (Spurrett and Cowley 2010, Cowley 2007).

This chapter explores and defends the extended mind thesis in detail. Doing so necessitates veering away somewhat from the largely phenomenology-driven considerations of previous chapters and adopting a more straightforwardly philosophy of cognitive science-based approach. However, in an attempt to warm the reader to the suggestion that their mind might not always stay inside their head, I begin by briefly revisiting some phenomenological insights concerning the nature of our mental lives found in Heidegger's early phenomenology. I then proceed to recount the arguments in favour of the extended mind thesis, focussing mainly on the

locus classicus of the debate, Clark and Chalmers' seminal *Analysis* article, 'The Extended Mind'. I then discuss several examples, both empirical and theoretical, of the extended mind thesis at work, and diffuse some commonsense objections. The bulk of the chapter tackles the most prominent criticisms of the extended mind thesis.

2. What it is like to have an extended mind

For many philosophers, cognitive scientists, and laypeople alike, the very idea that the mind might be partially located outside of the brain is so intuitively unpalatable as to be borderline intellectually offensive. Before entering into the specific arguments for and against the extended mind thesis, some preliminary phenomenological analysis might help to soften these prejudices somewhat.

How obvious is it, from the perspective of ordinary experience, that thinking - reasoning, planning actions, solving problems, etc - takes place inside our heads? Husserl for one, thought that this was not obvious at all

> [T]hinking is not actually located in the head, the way the impressions of tension are, etc. That we very often speak as if it were so is no proof that we actually apprehend it that way in intuition. (Husserl 1989 p.161)

Historically, speaking, it cannot always have been considered obvious. If it had, Aristotle would probably not have been able to entertain the hypothesis that the function of the brain was to cool blood, while the heart was centre of sensation and action.[40] Of course, it is certainly true that we often silently soliloquise when, for instance, trying to work out how best to verbalise a complex matter, or when mulling over a set of instructions. In these private cases it seems correct to observe, as

---

[40] See Gross (1995) for a detailed analysis of Aristotle's views on the heart and brain

Dainton does in his discussion of the spatiality of experience, that 'If I focus on my own conscious thinking, it apparently occurs somewhere within my head - between my ears, behind my eyes' (Dainton 2005 p.60). But not all thought is like this. For one thing, not all thinking is conscious, or deliberate. Moreover, a great deal of our ordinary day-to-day activities involve the application of practical reasoning, bodily skills, and interactions with our environment that while rarely the product of deliberative, conscious thought processes, are nevertheless intelligent, goal-directed behaviours with a distinctive experiential character.

Heidegger argued that it is our pre-reflective relationship with an already familiar environment, and not an attitude of detached intellectual scrutiny, that constitutes the default mode of being (or, if you prefer, state of mind[41]) for Dasein, his term for the sort of existence belonging essentially to humans. Prior to reflecting on the world around us, we typically find ourselves comfortably immersed in a network of familiar practices involving various tools and other assorted paraphernalia which we do not experience as foreign objects. In Heidegger's terminology, these tools are experienced as 'ready to hand', as opposed to 'present at hand' (Heidegger 1962 p.103).[42] For example, though it is trivial that my computer is a physical object located outside of my body and that in using it I perceive it, when immersed in the activity of researching current affairs on the internet or rectifying the shortcomings of my latest thesis chapter, I am not aware of it *as* a foreign object, or as something I am deliberately visually inspecting. Rather, my experience is directed at the activity in which I am engaged, and my consciousness is occupied primarily by the particular

---

[41] Heidegger would probably disapprove of this terminology, having gone to great lengths to avoid words like 'mind' for fear of Cartesian implications, but it is, I think, a sufficiently innocuous rendering of Heidegger's basic idea into commonsense language.

[42] For a thorough discussion of Heidegger on tool use, see B. Preston (1998).

project at hand. This harmonious engagement with the computer as the vehicle of my

current activity is only disrupted when something goes awry, for example if my

internet connection dies or (as is apparently commonplace with Toshiba latops!) it

spontaneously shuts itself down, or if I am adjusting to the use of a new piece of

software or hardware. The same goes, mutatis mutandis, when driving a car or

playing a musical instrument. Navigating roads or improvising using a pentatonic

scale are thoughtful activities in which decisions are made and actions performed on

their basis with a view to achieving a particular end, which are made possible by an

experiential unity with my surroundings in which the boundaries between the

thinking subject and the world in which they are situated are obscured. Heidegger

dubbed these familiar objects as they appear to us in ordinary experience 'equipment',

as illustrated in his well-known discussion of hammering

> Equipment can genuinely show itself only in dealings cut to its own measure
>
> (hammering with a hammer, for example); but in such dealings an entity of
>
> this kind is not *grasped* [i.e. understood] thematically as an occurring Thing,
>
> nor is the equipment-structure known as such even in the using [...] In
>
> dealings such as this, where something is put to use, our concern subordinates
>
> itself to the "in-order-to" which is constitutive for the equipment we are
>
> employing at the time; the less we just stare at the Hammer-thing, and the
>
> more we seize hold of it and use it, the more primordial to our relationship to
>
> it become, and the more unveiledly [sic] is it encountered as that which it is -
>
> as equipment...The kind of being which equipment possesses - [...] - we call
>
> *readiness to hand.* (Heidegger 1962 p.98, italics in original)

In our familiar dealings with the useful objects on which we rely to accomplish

various practical goals, we do not experience ourselves as isolated from our

environment, but deeply immersed in it, hence the strong opposition to the traditional dichotomy of subject and object than runs throughout *Being and Time*.

It is not just phenomenologists who are impressed by the way in which a great deal of our ordinary behaviour depends on our oneness with surrounding objects. The psychologist David Kirsh, whose research is motivated by scepticism regarding classical psychological models of action-planning whereby subjects must represent an abstract model of a task and the outcome of actions executed in order to complete it (e.g. Glover 2004), stresses the extent to which our mental lives and our capacities to exercise our intelligence are inextricably bound up with the environmental props we rely upon to get by in the world (I discuss some of Kirsh's empirical work and its relation to the extended mind theses below).

> We live most of our life in constructed environments. Layers of artefacts saturate almost every place we go, and there are pre-existing practices for doing things. These artefacts and practices have been designed, or have coevolved, to make us smarter, to make it easier for us to solve our problems and perform our habitual tasks. (Kirsh 2009 p.270)

Everywhere we go and in almost everything we do, we are surrounded by meaningful public items which play an important role in the governance of our behaviour, from road signs which aid us in remembering where to walk or drive, to items displayed in a supermarket in such a way as to affect our purchasing decisions, to mobile phones and tablets which store valuable information and run an ever-expanding array of labour-saving applications. On my understanding, the extended mind thesis, in essence, represents an attempt to situate these sorts of Heideggerian phenomenological insights within the context of contemporary cognitive science in a

way that is both empirically viable and compatible with the terminology of analytic philosophy of mind[43]. Of course, merely appealing to a carefully selected chunk of Heidegger's phenomenology does not cut the mustard as an argument for the extension of cognition into the world, but it does give the extended mind thesis some intuitive appeal by relating it back to our ordinary experience prior to its dissection in the lab or the philosophy classroom.

3. The parity principle

The extended mind thesis adheres to the same broad commitments to physicalism and functionalism common to the vast majority of contemporary philosophy of mind and cognitive science, but differs from mainstream cognitivism insofar as it rejects the widely held *additional* assumption that Wilson (2004) dubs 'smallism', namely the assumption that the realisation of mental phenomena, achieved by dint of the functional organisation of its composite physical processes, occurs only inside brains. Strictly speaking, this assumption of smallism is extrinsic to both physicalism and functionalism per se. As such, extended mind theorists take this assumption to be unjustified, overly dogmatic, and open to revision.

Central to understanding the extended mind thesis is the notion of *coupling*. An agent is said to be 'coupled' with an external object when reciprocal causal interaction between the agent and an object drives along some activity undertaken towards some particular cognitive end. Take one paradigmatic example (due originally, insofar as I can ascertain, to Wilson (1994)), to which I shall return at numerous points

---

[43] Clark often cites Heidegger as an influence, though refrains from discussing his phenomenology in any detail (e.g. Clark 2003; 1997). For a more through treatment of Heidegger's relevance to, and influence upon, contemporary cognitive science, see Wheeler (2007) and the essays collected in Kiversten and Wheeler (2012).

throughout the chapter, of a mathematician using a pen and paper to help solve an equation. The use of representational media outside of the body - diagrams, written word and symbols, rough sketches, etc - quickens the process of solving equations (Villegas, Castro & Gutierrez 2009). The cognitive end in this case is producing a solution to a given mathematical problem. The mathematician's act of writing causally impacts upon the paper by creating meaningful and useful marks on the page. These marks, in turn, causally impact upon the agent via visual perception. The activity of solving the equation on paper therefore comprises a two-way loop between an embodied, perceiving agent, and an external source of information. This loop exists so long as the mathematician is actively engaged in using pen and paper to perform the various computational operations required to reach their solution. The mathematician and their trusty stationary therefore comprise a particular kind of system, a *coupled system.* There is nothing remarkable or controversial about the idea that they instantiate a system - just about any set of coordinated components can be treated as a unified system (see, e.g. Mainzer 1997, Van Gelder 1998). What is arguably more controversial is Clark and Chalmers' stance that coupled systems like these are *cognitive* systems

In these cases, the human organism is linked with an external entity in a two-way interaction, creating a *coupled system* that can be seen as *a cognitive system in its own right.* All the components in the system play an active causal role, and they jointly govern behaviour in the same sort of way that [internal] cognition usually does. If we remove the external component the system's behavioural competence will drop, just as it would if we removed part of its brain. Our thesis is that this sort of coupled process counts equally well as a cognitive process, whether or not

it is wholly in the head. (Clark & Chalmers 1998 pp.8-9, first italics in original, second added)

This inclusion of extra-neural items into the cognitive domain is motivated by appeals to parity of reasoning between the inclusion of processes internal to an organism (i.e. in the brain) as part of its cognitive architecture, and items external to the organism which perform a coarsely isomorphic function. In other words, *where* an object is in relation to a person has no immediate bearing on whether it forms part of their mind. What matters is the role played by that object in underpinning their thoughts and behaviour. If an object external to the body plays the right kind of functional role, there is no non-dogmatic reason for discounting it as part of a cognitive state or process. This is the 'parity principle', due originally to Clark and Chalmers (though in their original article, they do not use this exact term), which Clark states thus

Parity Principle. If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is (for that time) part of the cognitive process. (Clark 2008 p.77)

It must be stressed that the extended mind theorist does not claim the functional isomorphism between an external process and its nearest internal counterpart must be especially fine-grained. As Clark explains,

[F]ar from *requiring* any deep similarity between inner and outer processes, the parity claim was specifically meant to *undermine* any tendency to think that the shape of the (present-day, human) inner process sets some bar on

> what ought to count as part of a genuinely cognitive process, (Clark 2012
>
> p.285, italics in original)

In what follows, I elaborate on a range of cases in which the parity principle instructs us to regard items outside of the brain and body, when suitably coupled to an intelligent agent, as part of their cognitive architecture, that is, as part of the physical basis of their mind. The conclusions which emerge from these considerations of parity are then defended at length against objections raised by more conservatively minded philosophers who wish to keep the mind firmly inside the head.

One pedantic note about terminology: the rhetoric of minds 'extending into' their surroundings is a little unfortunate. As one defender of extended cognition remarks, the image of a neural process literally travelling outside of somebody's head into an object suggests a metaphysics that is 'positively medieval' (Menary 2010b p.609). This isn't really what extended mind theorists mean when they say cognition is 'extended'. What they actually mean is that in instances of cognitive coupling, the processing involved is *distributed* across a surface area wider than that of select parts of the brain. If it were up to me, I would stick to more literally accurate talk of 'distributed cognition' but, alas, it isn't up to me. For better or worse the 'extended' label has stuck, while 'distributed cognition' has come to have a slightly different meaning in a lot of cognitive science, where it is generally used to refer to the distribution of labour in coordinated group behaviour (see, e.g. Maglio, Kandogan & Haber 2003, Hollan, Hutchins & Kirsh 2000).

4. From epistemic action to extended cognition

The ability to skilfully manipulate our immediate environment is of central importance to a great deal of our practical reasoning. As Kirsh notes

> [P]roblem solving is an interactive process in which subjects perceive, change, and create the cues, constraints, affordances, and larger-scale structures in the environment, such as diagrams, forms, scaffolds, and artefact ecologies that they work with as they make their way toward a solution. (Kirsh 2009 p.290)

According to the extended mind thesis, the use of such artefacts to simplify and quicken our problem-solving prowess extends the physical basis of users' minds by distributing the requisite cognitive processing across reciprocal loops between brain, body and world.

Clark and Chalmers appeal to the cognition involved in the retro videogame Tetris to illustrate this claim. In Tetris, puzzle blocks of various shapes (called 'zoids') fall from the top of a two dimensional grid. Points are scored by rotating and aligning the shapes to form complete rows which then disappear from the screen, freeing up space in the grid. Failure to make whole rows out of the blocks will result in the zoids piling higher and higher until the grid is full, at which point the player is, in the language of gaming, 'dead'. The cognitive task is therefore to identify the best shape of fit and available slots for the zoids. In a fascinating empirical study, Kirsh and Maglio (1994) found that players of all levels of ability tend not to move zoids in accordance with prior decisions concerning where they are to placed, but instead physically rotate them on the screen to assign prospective locations in real time. Skilled players frequently make moves that are superfluous from the point of view of traditional 'planning-control' models according to which puzzle solving is a matter

of internally representing a task space and computing potential solutions *and then* performing the corresponding action as a result. In actuality, players begin to manipulate zoids even before they come into full view on the screen (ibid. p.524). This suggests they do not proceed via a linear 'think-decide-act' strategy, but rather 'offload the structure [of the problem] onto the world', thus 'pre-empting the need for certain representations or…the need for making certain inferences' (bid. p. 545). Human subjects, even those with minimal gaming experience, easily outperform the control experiment –a sophisticated computer simulation of Tetris playing utilising an efficient linear, rule-following strategy (ibid. pp.518-26). This is because they do not face the temporal constraints imposed by the requirement of planning moves prior to execution. Neither do they face the additional computational effort incurred by generating a complete internal representation of the task space. Tetris players really do, to adapt Brooks's slogan-come-cliché, use the world as its own model (Brooks 1991a; 1991b).

The manipulation of the zoids to identify the best shape of fit is an example of what Kirsh and Maglio dub *epistemic actions*: 'physical actions that make mental computation easier, faster, or more reliable…that an agent performs to change his or her own computational state' (Kirsh & Maglio 1994 pp.513-4). Other examples of epistemic actions include moving chess pieces to discover the best available move (ibid p.545), and shuffling scrabble tiles to identify possible letter combinations (Kirsh 1995) (a fortiori in its mobile app equivalent, Words With Friends, where the entire tile try can be shuffled at the tap of a button). In these instances, part of the world is skilfully manipulated in order to visually identify the solution to a given problem, thereby alleviating the subjects' internal mental workload.

The psychological moral of epistemic actions is that 'It is both cognitively easier and computationally simpler to use the external representation than an internal representation' (Kirsh 2009 p. 279). The evolutionary moral is that 'In general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned' (Clark 1997 p.46). The *philosophical* moral, according to Clark and Chalmers, is that some of the cognitive processing needed to solve the puzzle literally takes place on the screen rather than exclusively in the players' brains, because the act of physically rotating shapes on a two-dimensional grid at the click of a button is coarsely functionally isomorphic to mentally rotating images 'in the mind's eye' to complete the same task (Clark & Chalmers 1998 pp.10-12; see also Clark 1997 pp.64-5; 203-4 and Clark 2008 pp.70-5). The flow of information and causal profiles of the two cases might not be entirely identical, but there is no reason to think that this should have any bearing on their status as part of a cognitive system; there are often myriad available internal procedures for completing the same cognitive task - think of the numerous associative techniques we can employ to aid memory recall, or the different formulae we could employ to solve the same mathematical problem. The only major differences between the way people actually play the game and the unrealistic alternative scenario of players mentally rotating imaginary shapes prior to hitting 'enter' are that (a) in the former case, the activity is distributed across the embodied agent and the computer rather than housed in their brain, and (b) access to the desired destination for the zoid is perceptual rather than introspective or imaginative (or whatever). (a) can only constitute the basis of an objection if it is assumed that the boundaries of brain and body constitute the bounds of the mind, but such an

assumption begs the question against the extended mind theorist. (b) is more of a fact about the associated phenomenology of different methods of Tetris playing than anything else, and it would be both overly strong and seemingly ad hoc to claim that two instances of cognitive processing must be accompanied by identical phenomenology in order to count as instances of the same type.

Another good example of cognitive extension via epistemic action, and one that is more common for most of us at least, is the role played by word processing in the writing process. When writing (say) a philosophy paper, we do not proceed by first privately formulating the arguments, clarifying the key ideas, and articulating these into readable English in our heads, and then typing them out by rote resulting in a finished article. Perhaps there are a tiny minority of technophobic savants for whom this would be a natural enough way to get the job done, but for the vast majority of people this would be practically impossible. Ideas are clarified as we go along, by reading, modifying and re-reading the wording and structure of the text, cutting, pasting and deleting sections as appropriate, crosschecking the text with other documents, etc. The use of external media is essential to the process of producing the article, not just in the trivial sense that without it there would be no way to record one's finished ideas, but because the writing is literally part of the thought process. The intellectual labour does not precede the writing process; the writing process is a significant and indispensible component of the intellectual work itself.

Though it is possible in principle to compose an academic article internally and subsequently record it on the page, typically we do not do this. We do not do this because it is inefficient and cumbersome, and because we simply do not need to,

because there exist malleable external structures which relieve our brains of a large amount of the cognitive workload. As Menary puts the point

> The media function as enabling hardware, but the vehicles themselves enable processes that cannot be completed in the head alone. The physical act of typing necessarily involves external physical manipulations. My ability to compose a paper is severely curtailed by the absence of those external manipulations [...] Stable and enduring external written sentences allow for manipulations, transformations, reorderings, comparisons and deletions of text that are not available to neural processes. (Menary 2007b p.629).

It is at best abstract and artificial, at worst a complete fiction, to think of the act of writing as enjoying a wholly autonomous existence from thinking. Word processing software makes available various actions essential to the process of writing - cutting, copying and pasting, highlighting, not to mention the ready visual display of text - that are not available to brains alone. The ability to formulate the content of a philosophy paper is so inextricably bound up with the usage of external media that the idea of focussing purely on what occurs in the writer's brain, vital and interesting though this is, seems to leave out something essential to an account of this particular species of intelligent activity. A cognitive scientist who did so would only be studying one isolated part of a broader explanandum. Furthermore, the example of essay writing satisfies the parity principle. Consider an internalist scenario where our would-be author attempts to privately compose the text in her head, commits it to memory, then goes about typing it up word for word. When formulating her argument and checking her wording, she would need to represent to herself whichever sentence(s) she was presently mulling over. The state of her brain which

encodes this representation would subsequently play a causal role in bringing about the act of typing up the final version of the sentence(s). How she does this exactly could vary. She might visualise the sentence to herself or privately rehearse it in a sort of internal dialogue, 'behind her eyes and between their ears', as Dainton puts it. How exactly she does it doesn't strictly matter; the important point is that this method of formulating a piece of text would involve representing the content of the article to herself in its various stages of completion. In the more common case of formulating the text using the computer display, something analogous and from the perspective of functional decomposition, coarsely isomorphic, happens. The only difference is that the work in progress is displayed for the author not in the form of a private quasi-visual or quasi-auditory (or whatever) image, but on the screen in front of them. Again, through epistemic action, much of the cognitive workload is distributed across brain, body and world.

5. Extended beliefs: the curious case of Otto and Inga

Thus far I have focussed on mental process, but what of mental states? By far the most discussed of Clark and Chalmers' original examples of extended cognition is their 'Otto's notebook' thought experiment. The thought experiment is designed to show that external objects coupled to an intelligent agent can form part of the physical realisation base for one or more of their dispositional beliefs. Clark and Chalmers invite us to imagine Otto, a man who suffers from a mild form of Alzheimer's disease and as such has trouble remembering certain pieces of information, namely addresses. To redress this deficiency, Otto uses a notebook to store addresses, which he keeps with him as much as possible. One day Otto desires to visit the Museum of Modern Art in his native New York, and in order to get there,

looks up the address in his notebook, which informs him that the museum is located on 53rd Street. Otto's counterpart in the thought experiment, Inga, also desires to visit the Museum of Modern Art. Inga is fortunate enough to have a healthy brain, which enables her to recall the location of the museum before setting out. Inga clearly has the belief that the museum is located on 53rd Street, and this belief is operative in guiding her behaviour, even though prior to her act of recollection this belief was neither conscious nor occurrent (Clark and Chalmers 1998 p.12). According to the extended mind thesis, Otto's notebook plays a functional role sufficiently similar to the neural encoding in Inga's brain which enables her to recall the location of the museum to count as an external component of his memory system, and it is equally correct to ascribe the dispositional belief 'The Museum of Modern Art is on 53rd Street' to both of them. Like Inga's biological memory, Otto's extended memory is reliable (the information therein is usually true), readily available (he keeps the notebook with him as much as possible), the information therein is endorsed automatically upon retrieval, and it guides his behaviour in much the same way, allowing us to reliably predict his behaviour via the attribution of a belief about the location of the museum. And while it is true that Otto's notebook is susceptible to damage, loss, or outside interference, these are also true, albeit to a lesser degree, of Inga's biological memory. Inga's memory can be similarly damaged or lost through injury or intoxication, and is open to manipulation from suggestion or physical interference at the hands of crafty neuroscientists (ibid. p.15). The only substantial difference is that Otto's access to the information about the museum is perceptual rather than introspective, but again, this difference is more phenomenological than anything else.

To many people the suggestion that Otto's notebook comprises part of his memory will sound crazy. Driven by this intuition, some philosophers have appealed to features of our commonsense folk psychological concept of belief to block the attribution of the belief that the Museum of Modern Art is on 53rd Street to Otto. As Chalmers points out, he and Clark are not especially interested in the finer details of whether or not their thought experiment is fully compatible with our ordinary folk concept of belief, remarking that the issue is merely terminological (Chalmers 2008 p.xiii). What matters, Chalmers says, is the extended mind thesis' potential for explaining the behaviour of coupled agents

> The deeper point is that extended states can function in explanation in very much the same way that beliefs function, and they should be regarded as sharing a deep and important explanatory kind with beliefs. This explanatory unification is the real underlying point of the extended mind thesis. (ibid. p.xiv)

Not being much of an epistemologist, I am sympathetic to this view. But it is hard not to read Chalmers' comments as something of a cop out. After all, Clark and Chalmers did choose to explicitly defend the claim that there can be extended beliefs. I shall therefore take it upon myself to defend the notion of extended beliefs against two prominent folk psychology-driven objections to extended beliefs, one due to Preston (2010) and another to Weiskopf (2008).

The immediately obvious alternative to Clark and Chalmers' characterisation of Otto's dealings with his notebook would be to say that Otto in fact has no belief about the location of the museum prior to consulting his notebook, and merely acquires it upon inspection, only to later forget it again. Preston peruses this more

conservative interpretation, citing Otto's apparent lack of first-person authority over the contents of his notebook prior to consulting it

> [U]pon being asked, Otto is the authority on whether what's written in his notebook is indeed what he believes. But he isn't authoritative about the contents of his notebook *before* he has consulted it. He can't avow what-he-believed-before-consulting -his-notebook at times prior to his consulting it, which is what matters to the claim that that he already believed it. So even though he isn't s*urprised* to find it out, Otto does have to wait and see how things are in the "external world" (i.e. his notebook) before finding out what he believes. However, in ordinary cases of belief [...], there's simply *no such thing as* "finding out what one believes." (The phrase can only mean *making up* one's mind.) (Preston 2010 p.360, italics in original)

Preston is right to point out that Otto's lack of surprise upon accessing the contents of his notebook is insufficient grounds for attributing an extant belief to him; I am often unsurprised by events I read about in the news, this doesn't mean I already knew about them. However, there are different senses in which one can be' unsurprised' by a piece of information. If, as Preston presents the case, Otto is unsurprised by the contents of his notebook simply because he is habituated to following its instructions and is aware of no information to the contrary, then he may just be in the same situation as my cynical news reader. But, as I interpret Clark and Chalmers, the mental life of Otto differs slightly from this. One important feature of the phenomenology of recollection is that the thing recalled is apprehended by the subject with a feeling of recognition or familiarity. Husserl called this feature of memory 're-presentation' - in recollection the thing recollected is given as a re-presentation of something previously presented in experience (see, e.g. Husserl

1991). When we recall things effortlessly, mere absence of doubt suffices for automatic endorsement of the belief in question. If somebody asks me which house number I live at, ceteris paribus, I shall reply 'number twenty-six' without hesitation. I will not need to consciously follow any procedure for retrieving the information that I live at number twenty-six. If I am presented with the trickier task of remembering, say, the name of a person I once met very briefly, I may find myself mentally 'searching' for their name until I reach an 'a-ha!' moment. My frustration is brought to an end by the realisation, '*that's* the guy!', and I now apprehend the person's name *as recollected*. On my understanding of Clark and Chalmers, Otto's memory recall is more like the second case. Otto, being impaired, often needs to make some extra effort to follow a particular procedure for recalling the address of the museum, namely looking it up in his notebook, to 'jog' his memory. Now Preston could retort that if this is the case, and Otto does indeed experience the information in the notebook as recollected, then he already had that belief anyway, and the notebook acts merely a prompt. But as ex hypothesi *there is no neural representation corresponding to the museum's location for the notebook to prompt*, if Otto does indeed have an extant belief about the whereabouts of the museum, this can only be because the relevant information is encoded in his notebook. If we conceive of the role of the notebook in Otto's mental life along these lines, then his belief is on a par with an internal belief that requires a bit of mental effort to access. And there is nothing remarkable about these kinds of beliefs.

It is worth noting that *even if* it *is* accurate to say that Otto 'finds out what he believes' about the location of the museum, this does not establish that he didn't already have that belief. Beliefs for which people do need to be presented with evidence in order to 'find out' that they have them are commonplace, as illustrated by

the prevalence of implicit biases. Saul (forthcoming) synthesises a compelling body of empirical data showing that many in the philosophy profession are disposed negatively towards women, which can have a severe impact on their consideration for academic posts, their treatment at conferences, and so on. Many of the men and women guilty of being so biased will avow, quite sincerely, that they do not harbour any ugly beliefs about the value or ability of female philosophers and their dealings with other women, partners or family members for instance, may well be different. But there seems to be no other way of making sense of the prevalence of negative bias against women than to conclude that on some largely unconscious level, many people hold sexist beliefs. Though it was a huge over-exaggeration of Ryle's to claim that people come to know what they believe in the same way they in which they know the beliefs of others - through observation, he was at least correct about *this* class of beliefs, where people really do need to be presented with the evidence of their own actions in order to appreciate the tacit motivations behind them (see Ryle 1949 pp.160-73, cf. Davidson 1984). So again, given the prevalence of these internal beliefs, only an internalistic double standard could motivate the denial of Otto's belief on the grounds that he needs to consult 'external evidence' to know what he believes.

Preston also worries that Otto's inability to affirm his belief about the museum's location prior to consulting his notebook prevents the extended mind theorist from ascribing that belief to him. It is also true that his belief about the museum's location is only operative in guiding his behaviour during or shortly after he has consulted the notebook. But this doesn't pose a problem for the extended mind theorist either, because many dispositional beliefs are like this. There are plenty of internal beliefs to which we can only ascent upon conscious recollection, and which play no role in

guiding behaviour at all prior to or shortly after conscious recall. We all have beliefs about pieces of trivia that have virtually no application to anything we do that require substantial conscious 'searching' to recall. I have plenty of beliefs about, for instance, events depicted in films I saw a long time ago and didn't particularly enjoy and lyrics to songs that were played on the radio when I was a child. These play absolutely no role in guiding any of my behaviour bar some very specific and incredibly rare and/or improbable scenarios, and require considerable 'jogging' to come to consciousness (I might, for example, have to sing the melody of the offending song to myself). Furthermore, *I cannot avow what I believe about these things prior to conscious recollection*, because of the relative difficulty involved in retrieving them. But it is obvious and undeniable that these count among my beliefs. By parity of reasoning, then, the fact that Otto can only avow his belief about the museum upon conscious recall, and that prior to recall there is no observable behaviour to be explained by the attribution of that belief, poses no problems for the extended mind theorist; such beliefs are commonplace.

A second worry about the status of extended beliefs concerns their integration with extant internal beliefs. Weiskopf (2008) worries that unlike ordinary cases of belief, modification to the contents of Otto's notebook are not automatically integrated with other relevant beliefs. If for example, I am told that the philosophy department has been moved to another part of the York campus, it is likely that I will cease to believe that it is located in the Sally Baldwin building, if I read that a neighbour has died I will probably cease to believe they will put their bins out for collection on Wednesday night, and so on. Weiskopf claims in that in 'ordinary quotidian cases' this process of belief revision is both automatic and unconscious (ibid. p.268; 269). The contents of Otto's notebook, however, are not automatically integrated with his

extant beliefs. If he is told that the Museum of Modern Art has been demolished, his notebook is not automatically amended accordingly. Similarly, if one page of his notebook contains the museum's address, while a later entry on another page informs him that it has burned down, the latter piece of information will not override the former. Weiskopf takes this to mean that his extended beliefs cannot really be beliefs at all

> Beliefs are, as I will say, normally informationally integrated with, and updated in concert with, other beliefs [...] But most externally located mental states do not share this feature. So [...] they cannot be beliefs'' (ibid. p.268).

For reasons highlighted by Kyselo and Walter (2011), this conclusion simply does not follow. Even if it is the case the belief integration is *normally* rapid and unconscious, this is not an essential feature of beliefs per se (ibid. p.254). There are many instances where we fail to connect two pieces of remembered information either rapidly or prior to conscious consideration. Kyselo and Walter's example is visiting the Museum of Modern Art's cafe. You believe that the museum is on 53rd street, and that the cafe does a good latte. While walking to the museum, you suddenly remember reading recently that it has been torn down. Only at that point do you infer that the museum is not the place to go for a decent latte. You were previously privy to two pieces of contradictory information: that the museum had been torn down and that the museum cafe served good coffee. These facts were encoded somewhere in your brain and you would have been able to access and affirm either if need be, but it took some time and perhaps some prompting for your belief set to be updated appropriately (ibid. p.251). I frequently have this experience after walking to the bank on a Saturday afternoon to find it closed, despite believing that it

is Saturday, that banks close early on Saturdays, and that the bank is the place to go to make a deposit. These are all things I believe, but their integration is far from rapid. Moreover, it is precisely the fact that I *didn't* bother to consciously think about whether or not I should walk to the bank that prevented them from being integrated in time to save myself a wasted journey. There is nothing remarkable about any of this either. So, the fact that more time is needed in order for Otto's extended beliefs to integrate with extant internal beliefs, or each other, and that he has to make a conscious effort to update his notebook in the light of new information, poses no problem at all for the extended mind theorist.

Weiskopf also claims that the possibility of inconsistencies between Otto's internal and external beliefs somehow poses a problem for the extended mind theorist, as it renders Otto irrational, which means that the attribution of his extended belief yields no useful behavioural predictions

> If we suppose these external states [the notebook entries] to constitute Otto's standing beliefs, then we must say that he both believes that the museum is on 53rd and that it has been demolished. Minds that are extended in Otto's sort of way seem especially prone to falling into this sort of inconsistency and irrationality. The source of such error is the absence of the normal consistency-maintaining doxastic mechanisms located in the functional architecture of believers. (Weiskopf 2008 p. 270)

Again, I fail to see the problem here. Many if not most (if not all) ordinary people believe all sorts of contradictory things; they believe that spiders are dangerous despite also believing that spiders cannot harm them, they believe that it is morally wrong to steal but that it is perfectly acceptable to illegally download music, they believe that God is benevolent, that violence is bad, but also that God is wrathful and

vengeful. These are beliefs they will sincerely affirm when prompted and which serve to guide their behaviour in different ways at different times. Some contextual knowledge may be required to predict what they will say and do, but they are believers nonetheless. And recall that Otto suffers from mild Alzheimer's disease; to expect his behaviour to be as rational as Inga's is to demand far too much. Occasional lapses in rationality and subsequent difficulties in predicting his actions are to be expected, but these have no bearing on the status of his extended beliefs as beliefs.

The shortcomings of these folk-psychological objections share a common diagnosis. Though they attack the notion of extended beliefs from different angles, both Preston and Weiskopf's arguments adopt the same strategy of identifying a feature lacked by Otto's extended belief, then proceeding in an attempt to use the absence of this feature to undermine its status as a belief, all the while failing to acknowledge that internal beliefs which also lack this feature are positively ubiquitous. Both rest on an over-intellectualised caricature of ordinary human behaviour to smuggle in an internalistic chauvinism through the back door.

6. Defending the extended mind thesis

Above I provided rebuttals to some objections aimed specifically at Clark and Chalmers' notion of extended beliefs. For the remainder of this chapter I take up the challenge of defending the extended mind thesis against the most influential arguments levelled against it more generally. By far the most prominent critics of the extended mind thesis are Fred Adams and Kenneth Aizawa (e.g. Adams and Aizawa 2007; 2001), so the bulk of the section will focus on their work, although I shall also discuss related objections due to Robert Rupert (2009), who argues for a more

conservative 'embedded' cognition, and Mark Sprevak (2009), who takes issue with Clark and Chalmers' appeal to the parity principle and their particularly liberal brand of functionalism.

7. Notebooks don't think!

Though it should hardly need stating, it is worth stressing that neither Clark nor any other defender of extended cognition has ever claimed that the external objects utilised in instances of cognitive coupling -notebooks, mobile phones, computers, etc- are themselves cognitive subjects. The very suggestion that, say, Otto's notebook itself remembers lists of dates and addresses, or that my computer screen engages in reasoning about the puzzle game I am playing is so utterly absurd and embarrassingly outlandish that one would expect even the most ardent critic of extended cognition to do their opponents the courtesy of refraining from attributing it to them. Nevertheless, Adams and Aizawa, in a particularly antagonistic opening passage, write

> Question: Why did the pencil think that $2 + 2 = 4$?
>
> Clark's answer: Because it was coupled to the mathematician.
>
> That about sums up what is wrong with Clark's extended mind hypothesis.
> (Adams & Aizawa 2010 p.67)

This is a straw man argument if ever there was one. The extended mind thesis does not have it, falsely, that the mathematician's pencil somehow thinks when it is used as part of a problem-solving process. Rather, it is the mathematician who thinks; it just so happens that the pencil and paper she uses to help her do so forms part of the

physical apparatus (which includes her brain, or parts of it at least) that makes her thinking possible. Clark himself responds to this by applying the underlying logic of Adams and Aizawa's remark to the brain itself. Paraphrasing (and parodying) Adams and Aizawa's pencil and paper example, he writes

> Question: Why did the V4 neuron think that there was a spiral pattern in the stimulus?
>
> Answer: because it was coupled to the monkey. (Clark 2008 p.86)

As Clark explains, the intuitive awkwardness of the answer owes to the absurdity of the assumption underlying the question, namely 'the idea that a V4 neuron (or even a group of V4 neurons or even a whole parietal lobe) might *itself* be some kind of self-contained locus of thinking' (Clark 2010b p.83, italics in original). Individual neurons do not think, but they do function as part of a larger system which does. The particular neuron or group of neurons implicated in the monkey's perception of the spiral image is merely one important part of a broader physical system underlying the cognition involved in the monkey's dealings with the stimulus; it is not itself a bearer of whatever mental state monkeys find themselves in when they look at spirals. Once this simple fact is acknowledged, Adams and Aizawa's flippant joke about the mathematicians pencil is stripped of its rhetorical force.

8. Adams and Aizawa's tripartite critique: a tripartite critique

It is easy to accept that pencils and neurons do not think.[44] But all unhelpful rhetoric to the contrary aside, a sceptic about extended cognition might still have a hard time jettisoning the intuition that there is clearly *some* important difference (or

---

[44] With apologies to panpsychists (e.g., Strawson 2006).

differences) between the neurons firing in the monkey's head (or indeed anyone's head) and the mathematician's pencil, regardless of whether or not it is coupled to an intelligent agent. Obviously, this is true; these are two very different things. The brain is a vastly complex organ, the pencil is a very basic tool. What matters is whether their many differences have negative ramifications for the extended mind thesis. In their much-cited 2001 *Philosophical Psychology* paper, its various follow ups, and their 2007 book-length treatment, Adams and Aizawa develop three related arguments that seek to do justice to this intuition. Below I outline and respond to each of these in turn. I shall call these the *science argument,* the *coupling-constitution argument*, and the *content argument.* I shall argue that the science argument is straightforwardly empirically false, that the success of the coupling-constitution argument depends directly upon the success of the content argument, and that the content argument fails, taking the coupling-constitution argument with it.

9. The science argument

Adams and Aizawa make numerous attempts to empirically discredit the extended mind thesis. They are happy to accept that extended cognition is possible in principle and that future scientific advances have the power to make it a reality; that the mind is in the head is not touted as a logical or conceptual truth. They are also happy to accept that tool use and action play an important role in shaping our mental lives and that these are of interest to the study of intelligent behaviour. They also agree with extended mind theorists that 'the skull does not constitute a theoretically significant boundary for cognitive science' (Adams & Aizawa 2001 p.46), because merely pointing to the fact that the brain is located in the head tells us nothing about the

169

nature of cognition. Rather, they claim that the arguments given by Clark and Chalmers in support of the extended mind thesis fail to overturn what they take to be an almost unanimous scientific consensus - the belief in the 'contingent empirical fact' that the mind is firmly inside the head

> Our view is simply that as a matter of contingent empirical fact, when one looks at [cognitve processes], such processes happen to occur almost exclusively with the brain. Insofar as we are intracranialists, we are what might be called "contingent intracranialists", rather than "necessary Intracranialists". (ibid. p.53)

Adams and Aizawa characterise their case for 'contingent intracranialism' as a defence of scientific orthodoxy against its radical opponents in the embodied, enactive and extended research programs in philosophy of mind and cognitive science (Adams & Aizawa 2007 p.9). At times, though, they are wont to invoke cognitive-scientific orthodoxy as itself constituting evidence against the extended mind thesis, in order to place the burden of proof squarely on the extended mind theorist (e.g. Adams & Aizawa 2010 p.74). As an argument, this will not do. Adams and Aizawa cannot appeal to orthodoxy in order to defend that orthodoxy on pain of circularity. One cannot respond to a series of arguments in favour of a philosophical position merely by stating that the burden of proof lies with its proponents; they will simply agree and redirect the respondent to whatever proofs they believe they have already provided. And in any case, nothing about the extended mind thesis contradicts any neurological findings, behavioural studies, or the experimental methodologies used therein. What is at issue is the philosophical gloss we should put on the findings of practicing cognitive scientists, whether, for example, Kirsh and Maglio's Tetris study should be construed in terms of contingent intracranialism or

extended cognition. Pragmatically speaking, scientists are quite justified in not caring too much about these finer philosophical details and getting on with the business of designing and conducting experiments in order to accumulate valuable data. Insofar as our choice of interpretative scheme for this data is a philosophical and not an empirical matter, received opinion amongst empirical researchers in the cognitive sciences is largely irrelevant. That said, it might be a little too rough and ready to sweep the science argument under the carpet on this basis. Fortunately for the extended mind thesis, though, Adams and Aizawa's science argument itself fails on empirical grounds.

Appeal to orthodoxy, granting that contingent intracranialism does in fact accurately represent this orthodoxy, fails as an argument because it is precisely this orthodoxy that Adams and Aizawa have taken it upon themselves to defend. But there is also a more nuanced variant of their science argument that potentially has more bite. Adams and Aizawa claim that unlike their more conservative contingent intracranialism, the seemingly more radical extended mind thesis yields no interesting or testable laws. For example, Adams and Aizawa cite psychophysical laws such as Weber's first law[45] and chunking effects[46] in memory recall tasks as examples of the sorts of lawful regularities we should expect from investigation into the nature of mental processes. It is true that Otto's use of his notebook is unlikely to conform to the chunking effects typically exhibited in human working memory. A single page of a notebook can easily contain more than the so-called 'magical number

---

[45] Weber's first law states that the just-noticeable differences between stimuli are proportional to the value of the stimuli, see Fancher and Rutherford (2012) pp.167-71.

[46] The ways in which bits of information are grouped into meaningful 'chunks' in working memory, aiding and constraining recall, see Miller (1956).

seven' pieces of information often touted as the upper limit of human working memory.[47] But this by itself provides no objection. It could only do so if Miller's law (or something suitably like it) were conceptually necessary for something to count as working memory which, of course, it isn't. If a person or creature were discovered whose working memory capacity far exceeded that of typical humans, we would not conclude on that basis alone that they weren't really remembering! So the fact that the behaviour of coupled systems might not be governed by the same psychological or psychophysical laws as their decoupled counterparts constitutes no objection to their existence or status as genuinely cognitive phenomena. However, Adams and Aizawa go a step further, expressing scepticism regarding the existence of *any* empirically interesting laws governing coupled systems,

> [I]n contrast to intracranial processes, transcranial processes are not likely to give rise to interesting scientific regularities. There are no laws governing humans and their tool use over and above the laws of intracranial human cognition and the laws of physical tools...What are the chances of there being interesting regularities that cover humans interacting with all these sorts of tools? Slim to none, we speculate. (Adams and Aizawa 2001 p.61)

As it is the business of science to uncover law-like regularities, the extended mind thesis will fall at the first hurdle of scientific credibility if this is true. But is it really the case that 'there are no laws governing humans and their tool use over and above the laws of intracranial human cognition and the laws of physical tools?'. The answer is an overwhelming 'no'. Clark's own counterexample comes from the field of human-computer interaction (HCI). HCI studies the emergent behaviour of humans

---

[47] According to George Miller's famous study, the number of informational chunks stored concurrently in human memory is limited to seven, plus or minus two (Miller 1956, cf. Baddeley 1992).

and computers construed as unified systems. These studies uncover plentiful causal laws over and above those found in psychology textbooks or the basic physical laws governing the machines themselves, and often also make recourse to modelling techniques prevalent in cognitive psychology (e.g. information storage, transformation and retrieval in memory (Clark 2008, p.96)). Classic HCI examples include Hutchins' much-discussed account of the levels of information processing crisscrossing between agents and their environment in the navigation of a naval ship (Hutchins 1996). With regards to tool use generally, there in fact are so many examples of empirical research which uncovers lawful regularities governing human-tool interaction that listing individual examples would not do justice to the extent of the field's success (for a detailed review, see Baber (2003)). On an even grander scale, the application of complex systems theory to group behaviour teaches us that all sorts of mathematical regularities and power laws emerge from behaviour of systems comprised of human agents and their surrounding environment (see Mainzer 2007, pp.311-66). Contra Adams and Aizawa, these laws are not reducible to those governing the internal neural resources of individual agents and their casual interactions with external props; they really are sui generis laws that provide reliable, testable predictions of the behaviour of emergent phenomena; the science argument is empirically false.

Of course, Adams and Aizawa could concede the existence of these and similar laws whilst continuing to object to extended cognition on the grounds that these laws are too specific to the particular task at hand and as such are unlikely to generalise to all instances of cognitive extension. This should not trouble the extended mind theorist either. It would be facetious to dismiss the findings of a psychological study dealing with one particular kind of performance task on the grounds that the regularities

discovered in the subjects' behaviour cannot be generalised to model their performance on a completely different sort of task. Given the vast and wildly heterogeneous range of tools utilised by human beings, this same should be expected to apply *a fortiori* in cases of cognitive extension via tool use. And in any case, there are commonalities to all successful instances of tool use to supplement our thought processes, such as increased speed and efficiency with regards to problem-solving.

The final strand of Adams and Aizawa's science argument attacks what they take to be an underlying commitment of Clark and Chalmers to a widely discredited philosophy of psychology: behaviourism. The methodological behaviourism espoused by pre-cognitivist psychologists like Watson (1913) and Skinner (1957) and echoed in philosophical circles by logical behaviourists, most notably Ryle (1949), fell out of fashion in the wake of a devastating critique due to Chomsky (1959) and the rise of contemporary cognitive psychology and computational neuroscience. Adams and Aizawa allege that Clark and Chalmers are implicitly committed to behaviourism, because they appeal to Otto's disposition to behave in certain ways, namely to walk to the Museum of Modern Art on 53rd Street, to support their claim that his notebook serves as an external memory source

> In fact, we find that radical [extended mind] theorists tend to rely on a behavioural -not to say behaviourist- conception of the cognitive...Cognitive scientists have generally rejected such behavioural conceptions of cognition, since they allow that gigantic look-up tables might count as cognitive models. (Adams & Aizawa 2001 p.47)

This attempt to tar extended mind theorists with the brush of behaviourism is unconvincing. Clark and Chalmers do not claim that Otto's notebook constitutes

external memory storage just because his observable behaviour mirrors that of Inga, but because the content of the notebook plays the same sort of causal explanatory role as the internal mental representation a good post-Chomskian cognitivist would posit in the explanation of Inga's behaviour:

> It is the way the information [in Otto's notebook] is poised to guide reasoning (such as conscious inferences that nonetheless result in no overt actions) that counts. This is not behaviourism but functionalism. It is systemic role that matters, not brute similarities in public behaviour (though the two are of course related). Perhaps Adams and Aizawa believe that functionalism just *is* a species of behaviourism. If so, we plead guilty to the charge but find it less than damning.(Clark 2010a p.52, italics in original)

In fact, as Clark and Chalmers see things, extended cognition follows from a very general commitment to functionalism, which is an integral part of the scientific orthodoxy Adams and Aizawa are so keen to defend. I address an objection to the extended mind thesis based on Clark and Chalmers' reliance on functionalism due to Sprevak (2009) below, but for now it suffices to note that Adams and Aizawa's charge of behaviourism is inaccurate. So much for the science argument.

10. The coupling-constitution argument

Adams and Aizawa's attempts to empirically discredit the extended mind thesis fail. Their other two arguments are of a more philosophical nature. I shall call the first of these the 'coupling-constitution argument'. The intended upshot of this section is that the coupling-constitution argument holds no weight when taken on its own, as it

presupposes the success of another argument, discussed in the following section, to even get off the ground.

The coupling-constitution argument attacks what Adams and Aizawa take to be fallacious reasoning on the part of the extended mind theorist. As they see things, when Clark and Chalmers appeal to complex couplings between agents and their environments to illustrate the case for extended cognition, they commit a 'coupling-constitution fallacy'. This is

> [A] tacit move from the observation that process X is in some way causally connected (coupled) to a cognitive process Y to the conclusion that X is part of the cognitive process Y. The pattern of reasoning here involves moving from the observation that process X is in some way causally connected (coupled) to a process Y of type ϕ to the conclusion that X is part of a process of type ϕ. (Adams & Aizawa 2009 p.81)

Exposing this alleged fallacy supposedly shows that the extended mind thesis confuses causal interactions between cognitive processes and their environment with extensions of those processes into that environment, and/or casual interactions between cognitive systems and external objects with a broadening of those systems themselves. Adams and Aizawa illustrate the coupling-constitution argument with examples of familiar domestic contraptions causally influencing their surroundings. Components of air conditioning units and CD players interact with surrounding air molecules, but we do not thereby infer that the affected air molecules are part of the process of sound production or air conditioning, or that air molecules are themselves

part of a sound system or air conditioning unit[48] (ibid. p.84). As an aside, I'm not entirely convinced that these examples show what Adams and Aizawa would like them to. Although not the sort of details one could reasonably expect to find listed among the item's components in the manufacturer's handbook, prima facie, it is no great stretch of the imagination to make sense of the idea of the surrounding air molecules, while not a proper part of the air-conditioning or sound-producing devices, being an essential part of the process of sound-production or air-conditioning. After all, if the machines were placed in a vacuum their individual components could continue to function just fine but accomplish nothing, so the suggestion that the processes of sound production or air conditioning require more than just the presence of a fully-functional air conditioning unit or home stereo is hardly elaborate; the extended mind theorist might even use the analogy to their own advantage if so inclined. What Adams and Aizawa need is a positive account of what 'constitution' actually amounts to. But as Ross and Ladyman (2010) point out, not only do Adams and Aizawa make no effort to provide such an account, there is no corresponding causal/constitutive distinction in contemporary physics for them to lean on[49].

Philosophical niceties about domestic appliances aside, there is at least one example in the literature of a sympathiser of Clark and Chalmers actually deploying this kind of reasoning. Noë, in a brief discussion of the extend mind thesis in the closing chapter of *Action in Perception,* characterises the extend mind thesis as the claim that 'the environment can *drive and so partially constitute* cognitive processes' (Noë

---

[48] I am well aware of the fact there is technically no such thing as an 'air molecule', but for present purposes this doesn't matter.

[49] Having almost no knowledge of contemporary physics, I defer to Ross and Ladyman's authority on this point.

2004 p.221, my italics). To this particular passage, Adams and Aizawa's contingent intracranialist would be quite justified in replying that 'a process P may actively interact with its environment, but this does not mean that P extends into its environment' (Adams and Aizawa 2001 p.56), thereby resisting the claim of cognitive extension. There may well be other instances of the fallacy in the extended mind literature. But even granting that instances of the coupling-constitution fallacy are present in some discussions of extended cognition, these are more likely due to sloppy wording on the part of the author than any deep-seated trouble with the case for the extended mind thesis itself. Clark and Chalmers' actual arguments in favour of the extended mind thesis rest on no such reasoning from causal interaction to constitution. Recall that the reason Clark and Chalmers offer for construing items such as Otto's notebook and the Tetris players' computer screens as part of an extended cognitive process is not the effect they have on their internal cognitive resources, but their role in the functional economy of the user. It just so happens that in these cases this is more widely distributed.

The crux of the issue comes down to this: what exactly is meant in this context by 'non-cognitive'? Or put positively: by virtue of what does something qualify as 'cognitive'? Adams and Aizawa have much to say on this issue, and I discuss their positive account below. However, the key point to note at this juncture is that any such account of the cognitive/non-cognitive distinction must serve as a vital premise in any variation of the coupling-constitution argument. Without such an account, any theoretical division of a physical process into cognitive and non-cognitive components will be arbitrary, with only unstated and undefended intuitions to guide it. The coupling-constitution argument presupposes a principle of demarcation for the cognitive and, taken on its own, is powerless against the extended mind thesis.

Adams and Aizawa's coupling-constitution argument therefore stands or falls with the success of their third and trickiest argument, the *content argument,* which I address in the next section of this chapter.

Before tackling Adams and Aizawa's content argument, my defence of the extended mind thesis is obliged to take into consideration a more sophisticated variant of the coupling-constitution argument due to Robert Rupert (2009). Rupert argues that the extended mind thesis should be rejected in favour of the more conservative hypothesis of *embedded* cognition. The hypothesis of embedded cognition shares with the extended mind thesis the innocuous claim that our mental lives are highly dependent upon complex interactions with various environmental props, but stops short of the claim that these external props are partially constitutive of cognition itself. Opting for more of a philosophy of science-based approach than Adams and Aizawa, Rupert invokes Occam's razor and argues that the extended cognitive systems and processes posited by extended mind theorists should be eliminated from the operative ontology of practicing cognitive scientist in favour of an exclusive focus on the persisting core of internal cognitive resources housed inside the organism, hence his characterisation of his view as 'organism-centred cognition' (ibid. p.47). The cornerstone of Rupert's approach is what he calls his 'systems-based criterion' for determining whether or not something is to count as part of cognition, which he states as follows

> [A] state or process is cognitive if and only if it is a state of, or process occurring in, mechanisms that are elements of the integrated set members of which contribute causally and distinctively to the production of cognitive phenomena. (ibid. p.35)

Its fatal flaw is its reliance on the phrase 'cognitive phenomena'. This phrase occurs throughout Rupert's critique of the extended mind thesis, but is never precisely defined bar some rough and ready allusions to the explananda of cognitive science. Taken by itself, this is not particularly problematic. However, if the systems-based criterion is intended as a principle of demarcation of the cognitive from the non-cognitive, and it is repeatedly and explicitly presented as such, then it is flatly circular. Of course, it can't be false that a cognitive system is a system that is responsible for the production of cognitive phenomena, but this is basically a tautology. Far from providing a principle of demarcation that might help work against the extended mind thesis, Rupert's systems-based criterion presupposes one. Hence my characterisation of Rupert's position as a variant of the coupling-constitution argument; Rupert wants to drive a wedge between cognition proper and media external to the organism by showing that it is a mistake to construe them as unified systems, but without being able to ground this in a clear-cut cognitive/non-cognitive distinction. Just like Adams and Aizawa's own coupling-constitution argument, Rupert's systems-based view requires but does not provide an account of what makes the cognitive 'cognitive' without which any division of coupled systems into cognitive and non-cognitive components will be arbitrary. In the next section I turn to Adams and Aizawa's attempts to do this and argue that they are unsuccessful.

11. The content argument and the 'mark of the cognitive'

The coupling-constitution argument is only serviceable if Adams and Aizawa can provide a principle of demarcation for the cognitive from the non-cognitive which, unlike Rupert's, does not suffer from circularity. They offer a candidate for such a

principle by means of what I shall call their 'content argument'. In this section I argue that Adams and Aizawa's content argument also fails, taking the coupling-constitution argument down with it.

Adams and Aizawa claim that as it is the business of the cognitive sciences to investigate cognitive phenomena, they (or philosophers acting on their behalf) owe us an account of precisely what it is for something to be 'cognitive' (Adams & Aizawa 2007 p.31). They call such a criterion the *mark of the cognitive.* They are keen to stress the search for a satisfactory mark of the cognitive should not descend into a mere terminological quibble, and offer instead what they take to be an empirical hypothesis about the nature of cognition based on the available evidence (ibid. p.49). Adam and Aizawa's candidate criterion is what they call 'intrinsic' or 'non-derived', as opposed to 'derived', content. Their distinction between derived and non-derived content is reminiscent of Searle's distinction between intrinsic and derived intentionality. Searle originally invoked the distinction within the context of machine intelligence (Searle 1980) before extending its application to debates surrounding the scientific understanding of consciousness (Searle 1992) and social ontology (Searle 1995). Searle famously argued (against the likes of Fodor (1980)) that the formal properties of a system - the syntactic structure of any program it physically implements - cannot suffice to make that system genuinely intelligent, because formal properties alone cannot guarantee any internal comprehension of the meanings of the symbols upon which the program operates. What Searle takes the system to require, is a prior understanding of the semantics of the program, and it is this 'original' or 'intrinsic' grasp of meaning that mere machines lack. As it is an undeniable phenomenological fact that humans do understand meaning and, on Searle's own understanding of contemporary neuroscience, an empirical fact that the

capacity to do so is somehow an accomplishment of the brain, Searle concludes that intrinsic intentionality is an as-yet unexplained biological property of humans', and possibly certain other animals', brains. It is this basic, biological endowment of human brains with intrinsic intentionality that supposedly explains the prevalence of socially constituted meaning; we are intentional agents who can grasp meanings, and as such can confer meaningfulness onto signs, flags, symbols, and the like in a secondary, derivative fashion (Searle 1983). Though they express scepticism with regards to his philosophically skewed reading of neuroscience (Adams and Aizawa 2007 pp.35-6), Adams and Aizawa are essentially in agreement with Searle over the distinction between intrinsic, and therefore genuinely mental, content, and its derivate counterpart in the public sphere

> Part of what we think distinguishes cognitive processes from non-cognitive processes is that cognitive processes involve representations. More specifically, cognitive processes involve non-derived, rather than derived, representations [...] Traffic lights, gas gauges, and flags are paradigm cases of items bearing derived content. Thoughts, experiences, and perceptions are paradigm cases of items bearing non-derived content. Roughly speaking the idea is that derived content arises from the way in which items are handled or treated by intentional agents. For the most part, things with derived content are assigned that content by intentional agents who already have thoughts with meaning. Underived content arises from conditions that do not require the independent or prior existence of other content, representations, or intentional agents. (ibid. p.32)

This bearing of non-derived representational content, as Adams and Aizawa see things, is a definitive feature of cognitive states and processes found only in brains,

which serves to demarcate the cognitive from the non-cognitive. If this is true, then items such as Otto's notebook, the mathematician's stationary, and the words and images displayed on the Tetris player's computer monitor, which contain only derived content, are ruled out as constitutive parts of their users' cognitive resources. The coupling-constitution argument, supported by a foundational premise to the effect that the cognitive is to be distinguished from the non cognitive on the basis of its exhibiting non-derived content, could be repackaged and redeployed, this time successfully, against the extended mind thesis. This is the *content argument.*

One possible way of responding to the content argument would be to offer an alternative mark of the cognitive that allows for extended cognition. Rowlands takes such an approach. Rowlands offers an alternative and extended cognition-friendly mark of the cognitive as information processing the proper function of which is to produce a representational state belonging to a subject (Rowlands 2010a pp.107-34). Though Rowlands' approach is interesting and I find myself in agreement with much of it, I do not think that this is a particularly fruitful avenue to pursue in combating the content argument. An Adams and Aizawa-style contingent intracranialist would be well within their rights to dismiss such a counter on the grounds that merely putting forward an alternative mark of the cognitive does not suffice to disprove or discredit their own, and is therefore a refusal to engage with their position rather than a direct counter to it. Furthermore, devising an account of the mark of the cognitive for the very purpose of saving the extended mind thesis runs the risk of begging the question against the contingent intracranialist.

Another tactic would be to deny outright any such distinction between non-derived and derived content. Dennett makes this move against Searle (Dennett 1996 pp.50-6), while Hutto and Myin follow suit against Adams and Aizawa (Hutto and Myin

2013 pp.135-53). Cards on the table: I am far more sympathetic to the positions of Dennett and Hutto and Myin than Adams and Aizawa's. That is, I am sceptical as to whether there really is a clear distinction between derived and non-derived content, let alone one of any relevance to cognitive science. And I find Adams and Aizawa's attempts to elaborate on their notion of non-derived content utterly perplexing. For instance, I am not at all convinced that 'we cannot make the firing of a particular set of neurons mean what it does, simply by agreement that it does (Adams & Aizawa 2010 p.72). This is precisely what happens when a new public symbol is coined, say, in the designing of a corporate logo. The marketers are in effect deciding what the neural state tokened by the perception of their logo shall henceforth represent. If their marketing campaign is successful, over time the logo will become synonymous with their brand and instilled in popular consciousness, and this will be what the neural states underpinning people's perceptions or mental images of the logo will now be about. The same goes, mutatis mutandis, for the teaching of any new words. So the very suggestion that 'the cognitive states in normal cognitive agents do not derive their meanings from conventions or social practices' (Adams & Aizawa 2001. p.48) seems to me to be patently false. Quite why we should need to (somehow) bracket out all the publically-derived content of people's mental states in order to explain their behaviour scientifically is completely beyond me, and I struggle in vain to discern what, if anything, the leftover 'underived' (sic) content of such states could even amount to. That said, I do not wish to delve into the issue of content fixation. A thorough treatment of this topic would take me too far afield and, in all honesty, I simply cannot understand what is going on in a lot of the relevant literature. I shall therefore follow a more conciliatory strategy in my reply to the content argument. I shall argue that even if a conceptual distinction between derived and non-derived

content is granted, it simply cannot do the work required of it by the contingent intracranialist.

There are at least three reasons why Adams and Aizawa's appeal to the derived/non-derived content distinction, even if there is such a distinction to be drawn, is doomed to fail. First of all, it is far from clear that cognitive scientists - real ones, not philosophers thereof - really do require anything like a mark of the cognitive. The history of science is littered with success stories that proceeded in either total ignorance or outright defiance of the supposedly a priori first principles dictated by philosophers. Biology, including its sub-discipline of artificial life (see, e.g. Gardner 1970), seems to get by just fine without getting bogged down in the philosophical complexities of formulating an adequate definition of 'life'. And the early modern conception of 'matter' and the atomism that comes with it- still influential in many philosophers' commonsense musings on physical systems - is embarrassed by contemporary physics (see e.g. the essays collected in Ross, Ladyman and Kincaid 2013). If these are to be the 'mature sciences' from which a 'would-be science' like cognitive science should take its lead (Adams & Aizawa 2001 p.53; 49), then the demand for a principled mark of the cognitive seems like an overly strong requirement bordering on pure philosophical fantasy.

Secondly, as Menary (2010b) points out, Adams and Aizawa themselves abstain from providing a positive account of how allegedly intrinsically content-bearing states get their content. They do, however, cite extant accounts from the naturalised semantics literature, including those of Dretske (1981), Millikan (1984) and Fodor (1990). Presumably, then, we are supposed to infer their theory of content from these. But these accounts of content are a heterogeneous and often competing bunch. Millikan's biosemantics, for example, is deeply rooted in Darwinism, while Fodor

advocates a strongly anti-Darwinian stance on the issue of representational content (Fodor 2008). Given the absence of a singular positive account of intrinsic content fixation, Adams and Aizawa's notion of intrinsic content is not so much a candidate mark of the cognitive as a purely theoretical stand-in for some as-yet unformulated mark of the cognitive to come by way of a fully developed and unanimously agreed upon naturalised semantics. But there is no such account, so quite what the extended mind theorist is supposed to respond *to* here is a complete mystery.

Thirdly and most importantly, even if we accept for the sake of argument that such a thing as the mark of the cognitive is required by cognitive science, and that the derived/non-derived content distinction can be adequately fleshed out, *and* that cognitive processes necessarily involve representations with non-derived content, this is still insufficient to equip Adams and Aizawa's contingent intracranialist with the right tools to deny cognitive extension, because Adams and Aizawa find themselves obliged to concede that not every state of a cognitive processes need exhibit intrinsic content anyway (!). In their original paper, they acknowledge the limitation of their own view by noting that

> Having argued that, in general there must be non-derived content in cognitive processes, it must be admitted that it is unclear to what extent every cognitive state of each cognitive process must involve non-derived content. (Adams and Aizawa 2001 p.50)

In their book-length treatment, they elaborate on this as follows

> As far as we know, cognitive processes might involve representations that include a small fixed set of non-representational functional elements, such as punctuation marks and parentheses. Such items might count as part of a

language of thought based on the manner in which they interact with items having non-derived content. If this happens, then cognitive states will to some extent be less than maximally dependent on non-derived content. *Not every part of every cognitive state will be content bearing*. (Adams and Aizawa 2007 p.55-6 my italics)

If even Adams and Aizawa are happy to admit that it would be too strong an empirical hypothesis to insist that every part of a process must bear non-derived content in order to qualify as a cognitive process, and that genuinely cognitive systems are highly likely to contain parts which do not bear such content, then an appeal to the presence or absence of non-derived content cannot be the deciding factor in whether or not it is correct to call a part of a process cognitive. When Clark (2010a) calls Adams and Aizawa on this, their direct response amounts to little more than a reiteration of the same fatal concession

> We don't think we have good reason to hypothesise that every component of every state of a cognitive process must bear non-derived content [...] Our view is that at least some components of cognitive states require some non-derived content, where *the states of a notepad for arithmetical computations, Otto's notebook, video games, and most mundane tools do not*. (ibid. p.50, my italics)

This harkens back to the fundamental misunderstanding of Adams and Aizawa's sarcastic joke about the mathematicians' pencil. The external item is not the bearer of a cognitive state - it plays a nontrivial role in the overall functional economy of a wider system which has that state. Otto and his notebook, Tetris players, mathematicians equipped with pencil and paper etc, are all systems which, ex

hypothesi, contain neural states bearing non-derived content, as well as processes or parts of process which do not. If the fact that these latter processes do not bear non-derived content does not automatically deem them non-cognitive after all, then Clark and Chalmers' arguments for the extended mind thesis remain untouched. As Clark himself puts it, at this point 'there is really no case (concerning intrinsic content) left to answer' (Clark 2010a p.49). The derived/non-derived content distinction, by Adams and Aizawa's own admission, fails to ground a clear-cut cognitive/non-cognitive distinction. The content argument therefore rests on an indefensible double standard in favour of the internal over the external. As such, it fails, taking the coupling-constitution argument with it.

12. The mark of the cognitive mk. II

Adams and Garrison (2013) offer an alternative candidate for a mark of the cognitive that shrinks away from the issue of content. They appeal to the role of reasoning in cognitive processing to outline what they take to be a necessary, but not sufficient, criterion, by which cognitive systems are to be distinguished from non-cognitive ones. In this section I flag up a number of concerns with their position and conclude that even if these are put aside, their mark of the cognitive sequel poses no real threat to the extended mind thesis either.

According to Adams and Garrison a system (an animal, a machine, or whatever) qualifies as genuinely cognitive if and only if its behaviour - by which they mean, somewhat confusingly, not just bodily movement, but the causing of that movement (ibid. p.341) - can be explained by reference to the system's reason(s) for behaving that way. There is a variety of senses in which we explain things in terms of reasons.

We might, for example, invoke gravity as the 'reason' for a precariously balanced object falling from a shelf, or a design flaw as the 'reason' for a glitch in computer software. Similarly, as Adams and Garrison remark, we might invoke facts about natural selection as the 'reasons' for the presence of a particular trait in a species. This use of the term 'reason' means something more akin to 'cause' or 'contributing historical factor', than 'reason' as it occurs in the sorts of explanations Adams and Garrison have in mind: 'cognitive explanations', that is, explanations of intelligent behaviour. They therefore distinguish their intended sense of 'reasons' from other common uses of the term

> [O]n our view of reasons that are involved in cognitive systems with cognitive processes, the reasons required must be the system's own reasons, and the explanation in terms of reasons must not only be at the evolutionary level of the species or programmer (or whatever). Call these *system centred reasons.* The explanations in terms of the reasons are teleological. Cognitive systems do one thing A in order to achieve another thing B. In the types of explanations we are featuring, the goal of doing B and the strategy for accomplishing B by doing A are represented within the system. (bid. p.347, italics in original)

Note that nothing in their definition of system centred reasons states that the system in question must be conscious of a goal state or their strategy for achieving it, or that these inner representations need even be accessible to consciousness. Adams and Garrison are quite happy that cognitive systems need not have conscious access to their own system centred reasons (ibid. footnote 24), and the strategy for achieving a particular goal might be encoded in an unconscious motor signal. Hence, what they have in mind with their notion of a system centred reason is not so much a *rationale*

that retrospectively justifies an action (in the manner of Davidson (1963)), but rather a sub-personal representation of the intended outcome of an action that causes that action. Genuine behaviour, in their stipulated sense, is behaviour that can be explained as the effect of an internally represented goal state

> [W]hether the motions of a system (or creature) have cognitive causes cannot be determined by observing the motions alone. For there to be genuine intentional behaviour explicable in cognitive terms depends upon the causes of the motions (movements). Only motions with cognitive explanations are truly cases of intelligent behaviour. (Adams and Garrison p.342)

I do not wish to quibble over their slightly unusual use of the word 'behaviour' or get bogged down in tangential concerns in the philosophy of action concerning whether or not reasons are causes (see the essays collected in D'oro and Sandis (2013)); Adams and Garrison may stipulate terms howsoever they wish. So for the remainder of this section, to avoid the appearance of begging the question, I shall reserve the term 'behaviour' for bodily movement causally explicable in terms of reasons.

Adams and Garrison invoke two examples from ethology and one from robotics. They claim that it is by virtue of a cat's representing a particular goal-state (e.g. the acquisition of food) and an accompanying strategy for achieving it (i.e. a particular sequence of bodily movements) that the stalking of its prey constitutes the work of a genuinely cognitive system, while the co-ordinated activity of a termite mound does not (Adams and Garrison take the termite example from Turner (forthcoming)). The cat has, represented somewhere within it, both a desire for food (the goal state) that is the cause of the bodily movement that brings about the satisfaction of the desire, and the motor program(s) needed to initiate these bodily movements. The 'behaviour' of

the termite mound, on the other hand, contains no such representations and can be explained in its entirety (i.e. using biochemistry) without the need for positing any such representations; the cat engages in reasoning, the termite mound, considered as a 'superindividual', as Turner calls it, does not.

Adams and Garrison also appeal to Rodney Brooks's robotic 'creature', Herbert, to support their mark of the cognitive. As noted very briefly in chapter I, Herbert is able to perform the task of collecting empty drinks cans dispersed around the floors of the MIT robotics laboratory whilst avoiding obstacles in the absence of any pre-programmed knowledge of how to recognise drinks cans or instructions on where to find them[50]. Herbert is a classic and oft-cited example of basic embodied intelligence in action, as it (he?) utilises neither a classical symbolic computational architecture nor, as Brooks is at pains to emphasise, a neural network (Brooks 1991a.). Unlike robots spawned by traditional 'top-down' Artificial Intelligence (Nilsson and colleagues' Shakey, for example[51]), Herbert's can-collecting emerges from the dynamic interactions between Herbert's active movement as conditioned by its bodily morphology (it is mobile and fitted with a grasping device) and its local environment (Brooks 1991a; 1991b). With no central processing unit with global control of its movement, Herbert's activity is entirely modulated by feedback signals distributed throughout its body. As Brooks explains

> The point of Herbert is two-fold. 1. It demonstrates complex, apparently goal directed and intentional, behaviour in a system which has no long term internal state and no internal communication. 2. It is very easy for an observer

---

[50]A video demonstration of Herbert can be found here https://www.youtube.com/watch?v=YtNKuwiVYm0 (last accessed 22/8/14).

[51] A video demonstration of Shakey can be found here: https://www.youtube.com/watch?v=RhrLHkVuerc (last accessed 22/8/14).

of a system to attribute more complex internal structure than really exists. Herbert appeared to be doing things like path planning and map building, even though it was not. (Brooks 1991a p.19)

In a move echoing Descartes' principled denial of intelligence to machines in his *Discourse on the Method* on the grounds that they lack reason (Descartes 1998), Adams and Garrison write

Suppose that a human lab assistant in the A.I. lab picks up a soda can and returns it to a trashcan. In some sense of behaviour, Herbert and the lab assistant have done the "same thing". The overt movements are the same. However, on the notion of behaviour that involves causing, these may not be the same behaviour. The lab assistant retrieves the can *for a reason.* Perhaps the reason is to remove clutter in the lab. Herbert's movements, however, are not performed for Herbert's reasons. Herbert has no [system centred] reasons. (ibid. p.342)

Supposedly, then, it is Herbert's apparent lack of system-centred reasons that renders it a non-cognitive system, and it is for this reason that Adams and Garrison rebuke as 'behaviourist' Brooke's characterisation of Herbert's activity as 'intelligence without reason' (ibid. footnote 14). Indeed, for them, 'intelligence without reason' is an oxymoron, and Herbert's 'behaviour' is not genuine behaviour, as they have defined it, at all.

There are notable lacunae in Adams and Garrison's own reasoning regarding these three examples. Insofar as Adams and Garrison deem Herbert's activity unintelligent and non-cognitive on the grounds that Herbert lacks reason *as they have defined it -* as having a representation of a goal state and a strategy for achieving it which causes

its bodily movement - they simply assume their candidate for a mark of the cognitive without arguing for it. The same goes for their treatment of the stalking cat and Turner's 'superindividual' termite mound. Adams and Garrison alternate between appealing to the role of reasoning in the cat's behaviour to justify the claim that the cat, but not the termite mound, comprises a cognitive system, and appealing to the cat's behaviour to support their hypothesis that cognition necessitates the presence of system centred reasons. These arguments are not only circular but secure a sleight of hand by appealing to popular and prima facie plausible intuitions - that cats have minds, that termite mounds don't, and that robots don't really know what coke cans are - then cashing these out in terms of a conception of 'reasoning' that is repeatedly stated but never explicitly defended.

As Adams and Garrison see things, the positing of system centred reasons is necessitated by the type of explanation proper to the understanding of intelligent behaviour, but no justification is given for this beyond the banal fact that the explanations of cognitive science explain cognitive phenomena and other types of explanation do not. They simply assume that it is the necessity of their system-centred reasons that grounds this difference: 'We maintain that if the full and accurate explanation of a system's behaviour must include the system's own reasons, then one is explaining the behaviour of a cognitive system and those reasons will be system centred reasons.' (ibid. p.347-8). What is required to make this work is a demonstration of the alleged fact that all cognitive processing necessarily requires representations of reasons, but no such demonstration is forthcoming.

In the absence of such a demonstration, let us be charitable and take it that the claim enjoys some prima facie acceptability. After all, the suggestion that cognitive systems are systems capable of behaving for their own reasons hardly sounds

implausible. The above reservations about Adams and Garrison's circular argumentative strategy notwithstanding, however, it is far from clear that this second mark of the cognitive actually can carve the sorts of distinctions they would like. While the issues of animal minds, collective mentality and machine intelligence are the subjects of ongoing controversy (see, e.g. the essays collected in Lurz (2009), Chant, Hindriks and Preyer (forthcoming) and Preston and Bishop (2002) respectively), I can agree with Adams and Garrison that cats have minds and termite mounds don't both for the sake of argument and because this is fairly uncontroversial. I will remain agnostic as to whether a robot like Herbert can justifiably be deemed 'intelligent', but I strongly disagree that Herbert lacks system centred reasoning *as Adam and Garrison have defined this*. This is highly problematic for their proposed mark of the cognitive if it is supposed to rule out Herbert while allowing for humans and cats.

To see why their example of Herbert fails, consider how Adams and Garrison might respond to the following potential line of objection. It could be argued that there is nothing immediately objectionable about the idea of the lab assistant picking up the can without having any particular reason to do so; she might be acting purely on impulse (perhaps she has obsessive compulsive disorder) or blind habit (her strict upbringing has conditioned her to be neat and tidy). When asked why she picked up the can, she might honestly reply, 'I have no idea'. Given the fact that her act of can-collecting involved the co-ordination of bodily movement in response to perceptual stimulation, it would be highly counterintuitive to regard it as the work of a non-cognitive system. This would constitute a case of genuine cognition which does not involve system centred reasoning, thereby disproving the claim that the presence of system centred reasoning is *necessary* for cognition (though it may still be

sufficient). Now, Adams and Garrison can respond that although the lab assistant might lack conscious access to her reason for picking up the can, this does not mean that she lacks a reason for doing so - in their sense of 'reason' - as not only can a system-centred reason be absent from the system's current conscious awareness, it can be inaccessible to consciousness completely. There is therefore nothing problematic in saying that the lab assistant could pick up a can either entirely absentmindedly (as it were) or while consciously thinking about something else, while continuing to ascribe to her a sub-personal representational state which causes and guides her can-collecting behaviour. The problem for Adams and Garrison is that if *this* is all having a system-centred reason amounts to, then Herbert *does* satisfy their mark of the cognitive. And if Herbert satisfies their allegedly necessary condition, then whatever morals are supposed to follow from Herbet's non-cognitive status for the extended mind thesis are mooted.

It is generally unwise to defer to a robotocist in philosophical matters, but it is even less advisable to overrule them when it comes to the empirical details of their own designs. Despite employing the rhetoric of 'intelligence without reason', and 'intelligence without representation' to describe Herbert (the titles of his 1991 papers), in his more sober moments Brooks stresses that Herbert is actually *not* an entirely non-representational device

> My earlier paper [Brooks 1991a] is often criticized for advocating absolutely no representation of the world within a behaviour-based robot. This criticism is invalid. I make it clear in the paper that I reject traditional Artificial Intelligence representation schemes (...). I also made it clear that I reject *explicit* representations of goals within the machine. There can, however, be representations which are partial models of the world - in fact I mentioned

that individual layers extract only those aspects of the world which they find relevant - projections of a representation into a simple subspace" [...]. The form these representations take, within the context of the computational model we are using, will depend on the particular task those representations are to be used for. (Brooks 1991b p.20, my italics)

What Herbert lacks is *not* the capacity to represent its surroundings (a la Turner's termite mound), but a *complete* internal model of its environment encoded in a central unit, and pre-programmed instructions about how to locate and collect discarded cans (Brooks's 'explicit goal representations'). Herbert does, however, represent *some* aspects of its environment which serve as cues for the particular movements it makes (i.e. avoiding approaching people by moving backwards, activating the grasping mechanism in the presence of something suitably coke can-shaped, etc). These minimal representations are distributed throughout Herbert's circuitry rather than housed in a CPU, but there is no non-chauvinistic reason for supposing that their location is of any importance. Recall from chapter I that Clark (1997) helpfully dubs these minimal, context-sensitive, action-guiding representational states 'action-oriented representations', which are 'representations that simultaneously describe aspects of the world and prescribe possible actions, and are poised between pure control structures and passive representations of external reality' (Clark 1997 p.49, cf. Wheeler 2005 pp.195-200). Why does this matter? Because Herbert's action-orientated representations satisfy Adams and Garrisons definition of a system-centred reason. Herbert moves around and initiates the particular movements it does because the required motor commands are triggered by states which represent certain features of his local environment. So while Herbert may lack an explicit representation of the overarching goal of collecting cans, it *does*

utilise representations of simpler, action-specific goals constructed on the fly. Of course Herbert isn't conscious of any reason for acting, Herbert isn't conscious of anything. But, ex hypothesi, this is irrelevant. To be clear, I am *not* claiming that Herbert is genuinely intelligent, or even that it deserves to be called a cognitive system. I remain agnostic on both fronts. The intended moral here is simply that Adams and Garrison's criterion doesn't even work in the context of their own example.

Thus far I have argued that Adams and Garrison's revised mark of the cognitive is vaguely formulated, poorly argued for, and unequipped to do the work they take to be required of it. But the crucial question is whether this mark has any bearing on the extended mind thesis, as Adams and Garrison maintain. I conclude this section by arguing that even if they are correct in holding that cognitive systems necessarily have system-centred reasons, this causes no problems at all for the extended mind theorist

Let us accept for the sake of argument that genuinely cognitive systems, such as those found in humans, cats, and whatever sorts of possible technologies Adams and Garrison are happy to allow, must of necessity be capable of behaviour in their stipulated sense of bodily movements caused by system centred reasons. Furthermore, let us be charitable yet again and suppose that all the issues raised about their method of argument and their conclusions regarding Brooks's robots can be adequately resolved. Would wholesale acceptance of Adams and Garrison's account in any way exclude Clark and Chalmers' coupled systems from qualifying as genuinely cognitive? The answer has to be a resounding 'No'. This is because Adams and Garrison's revised mark of the cognitive has absolutely no bearing on extended cognition whatsoever *unless* internalism is presupposed.

Consider our staple examples of Otto and his notebook and Kirsh and Maglio's videogamers. Both are clearly involved in forms of reasoning. Otto is planning a route to the Museum of Modern Art, and the videogamers are solving geometrical puzzles. Furthermore, both satisfy the more technically defined notion of a system centred reason. Otto's behaviour - his walking to the museum - is caused by a representation of his desired goal state, namely the location of the museum as depicted (in words, diagrams, or whatever) in his notebook. It just so happens that the goal state is represented in an external medium. The motor commands required for perambulation are caused by his processing, in this case via visual perception, of the representations in the notebook. Otto and his notebook, construed as a single coupled system, therefore satisfies Adams and Garrison's necessary criterion and qualifies as a cognitive system. The criterion can only serve to exclude Otto and his notebook if it is assumed that representations of goal states must be internal to the brain, which would beg the question, as well as ruling out from the armchair the sorts of possible mind-extending technologies they wish to allow for (Adams & Garrison 2013 footnote 22). Note that the extended mind theorist can even comply with Adams and Garrison's characterisation of the represented goal state as 'internal' to the system. The goal state is indeed 'internal', it is just 'inside' a more physically distributed system.

Similar morals apply, mutatis mutandis, to the Tetris example. Recall that the key moral of Kirsh and Maglio's Tetris study is that players use epistemic actions to manipulate the shapes displayed on the screen themselves in order to identify the best shape of fit for falling puzzle pieces *without* relying on an internally-generated model of the puzzle and the outcomes of possible moves. Once we have rejected the classical picture of detailed internal models of the puzzle pieces in favour of active

agents manipulating the task space itself, the objects comprising the task space take over the functional role of some internal goal states. Recall that Kirsh and Maglio found that players *do not process the outcomes of specific moves before making them*; outcomes of particular moves are discovered via the skilful manipulation of the shapes rather than being calculated in advance. If there are no prior decisions regarding the outcomes of specific moves, then are no goal states of *this* kind to account for. Nevertheless, there is a representation of puzzle pieces and an indication of where they will land given their current alignment and orientation: it's just not in the players' heads, but on the screen. This external representation is the cause of the player's act of hitting 'enter'. Again, this satisfies Adams and Garrison's definition of a system centred reason. Only internalistic prejudices, or recourse to further necessary conditions like a strengthened version of Adams and Aizawa's dead-in-the-water content argument, or some other and as-yet unoffered criterion, could do the work required to rule out extended systems while preserving wholly internal ones. When subjects enhance their reasoning processes via complex couplings with environmental structures, the makeshift systems which emerge still satisfy Adams and Garrison's allegedly 'necessary' criterion for cognition. Their mark of the cognitive: mark two is as impotent against the extended mind thesis as Adams and Aizawa's original.

13. Mars Aattacks!

Attempts to undermine the extended mind thesis by appeal to a mark of the cognitive fail even under the most charitable and conciliatory of readings. A novel line of objection along very different lines is pursued by Mark Sprevak (2009), who takes

issue with Clark and Chalmers' liberal version of functionalism. Recall that according to the parity principle, only coarse-grained isomorphisms between internal brain-bound cognitive states and processes and their externally distributed counterparts need hold in order for the latter to qualify as genuinely cognitive, a point Clark is at pains to emphasise (Clark 2008 p.88). To bolster the case for their argument from parity, extended mind theorists often appeal to possible beings - hitherto undiscovered aliens or animals, futuristic cybernetic organisms, and the like - whose internal cognitive machinery might well be very different from our own, but which enabled them nevertheless to entertain mental states and processes sufficiently similar to ours. The intuition underling such examples, which depending on the way in which it is deployed is either a fortunate consequence or a motivating factor behind functionalism about cognition, is that creatures could exist who are capable of entertaining the same sorts of mental states and processes as us despite these being implemented via some very different hardware (see Chalmers (1993) for a thorough discussion of implementation and multiple realisability). Sprevak dubs this the 'Martian intuition' (Sprevak 2009 p.8). In the context of the debate over extended cognition, the Martian intuition is often deployed by extended mind theorists in the following sort of argument: given that we would be happy to grant that these enjoy cognitive status, there is no non-chauvinistic to deny that status to an extended human process that works in the same way. Sprevak takes issue with this use of the Martian intuition[52]. He argues that the particularly liberal brand of functionalism at play in the argument from parity is too permissive, and that this undermines the use

---

[52] Sprevak ultimately concludes with an expression of scepticism not just over the extended mind thesis, but functionalism about the mind generally. Addressing this far stronger and more 'metaphysical' conclusion would take me too far afield, but for a discussion and response, see Drayson (2010).

of the Martian intuition to support the extended mind thesis. This is supposedly because we can easily imagine many forms of alien (or cybernetic, or whatever) cognition which when recast in terms of an analogous pairing of human agents with external props, in accordance with Clark and Chalmers' insistence that only a course-grained isomorphism between the two is sufficient to secure parity between the two cases, radically violate our intuitions about what deserves to be called cognitive and what does not.

Sprevak uses the Martian intuition to provide a reductio ad absurdum against the extended mind thesis. He suggests that the Martian intuition, combined with the parity principle, yields the extremely undesirable conclusion that simply by picking up a book, I come to believe everything in it:

> The justification is as follows. Consider a Martian [...] who encodes memories using ink marks. As well as acquiring beliefs via its senses, it seems possible for such a Martian to be born with innate beliefs. Furthermore, it seems possible for an organism to have innate beliefs that it has not examined yet - a library of data that is hard-wired into the organism by developmental processes, which the organism has not yet had cause to employ. Imagine that [...] [t]he Martian has ink marks inside its head that, if it were significantly diligent, would guide its action in appropriate ways; I have ink marks just outside my head that, if I am significantly diligent, would guide my actions in appropriate ways. The difference [...] is that it has the

ink-marks inside its head, while I have ink-marks outside. By the [parity

principle]53, if the Martian has the beliefs, then so do I. (ibid. p.21)

Given that, ex hypothesi for the extended mind theorist, spatial proximity is irrelevant to an external item's being part of a cognitive system so long as it has the requisite functional poise, the scope of the reductio is even more damningly broad. If this Martian example is accepted, then the whole university library could potentially count as part of my mind (ibid. p.22-3)!

If Sprevak's argument is successful, then the extended mind theorist is in deep water. Thankfully, however, a satisfactory rebuttal comes courtesy of Wheeler (2010). Wheeler points out that Sprevak's reductio rests heavily on the assumption that the Martian's innate and unexamined internal ink marks do indeed constitute a cognitive resource. But what reason is there for believing this? None is offered. In fact, as Wheeler highlights, what Sprevak does is to imagine particular examples which nobody, extended mind convert or not, would admit as an instance of cognition, transport it into the head of his Martian, then attribute the error to the extended mind theorist. As Wheeler himself puts it

> What Sprevak does [...] is take what he assumes to be the noncognitive, externally located elements in a distributed process, place them inside the head of a Martian, and conclude that they now deserve to be rewarded with cognitive status. But where is the justification for suddenly counting these elements as themselves cognitive? Apart from their spatial location, nothing

---

53 Sprevak himself uses the term 'the fair-treatment principle', but I have reverted to Clark and Chalmers' original terminology for the sake of consistency.

about them has changed from when they were judged to be noncognitive. The only new factor is their recently acquired-in-the-head-ness. (ibid. p.263)

Ironically, the only way such an argumentative strategy could work would be if it was assumed that a process counted as part and parcel of cognition purely by dint of its occurring inside a creature's head - that 'in-the-head-ness' is the mark of the cognitive - a claim which even Adams and Aizawa reject as ridiculous. Hence, 'Sprevak's argument depends on a form of the Martian intuition that is significantly more radical than the one he explicitly formulates as part of his conceptual backdrop' (ibid. italics removed). Once it is acknowledged that no one in their right mind, extended or otherwise, would admit mere vicinity in relation to a library as an instance of extended cognition, Sprevak's argument from the Martian intuition collapses.

Recall that Clark and Chalmers' use of the parity principle begins with the identification of a paradigmatic example of a cognitive process, such as remembering the location of a building or solving a geometric puzzle, then moves a part of the process outside the boundary of skin and skull. A question remains as to what, exactly, justifies the identification of some sorts of process - recalling information, spatial reasoning - as cognitive, but not others (and we are no longer accepting non-derived content as an answer!). Here Clark and Chalmers' rely heavily on commonsense intuitions. Making these intuitions explicit and defensible requires more philosophical heavy lifting. But unless Sprevak or someone acting on his behalf is prepared to produce a more nuanced account of these which justifies the inclusion of proximity to a bookshelf *by itself* on the same list as (inter alia) perceiving, believing and remembering, extended mind theorists can continue to hold that Otto believes the Museum of Modern Art is on 53rd Street, but not the entire contents of

the New York public library, without being guilty of, as Sprevak puts it, 'behaving like a NIMBY'[54] (Sprevak 2009 p.14).

14. Conclusion

This chapter explored and defended the extended mind thesis - the thesis that cognition does not take place in brains alone, but oftentimes is distributed via cognitive processing which loops across brain, body and world. Following Clark and Chalmers, I argued that cognition extends beyond the boundaries of skin and skull when embodied agents skilfully offload some of the processing that would otherwise have to take place internally onto external structures. A mix of empirical and theoretical examples were provided in support of extended cognitive processing during problem solving, and also of extended cognitive states, namely beliefs partly constituted by extra-neural objects. I addressed a wide range of objections to the extended mind thesis and found all of them wanting. Preston and Weiskopf's objections to the notion of extended beliefs fail because they aim to undermine extended beliefs by highlighting features they lack, but fail to appreciate that there are common instances of internal beliefs which also lack these features. Adams and Aizawa's tripartite critique of the extended mind thesis was also unsuccessful. Their science argument, it transpires, is straightforwardly empirically false. Their coupling-constitution argument accomplishes nothing on its own, as it presupposes a principle of demarcation it does not provide. Rupert's systems-based critique fails for the same reason. Adams and Aizawa's attempt to provide such a criterion, their content argument, does not provide the requisite conceptual resources to resist extended

---

[54] 'Not in my back yard'.

cognition, even if the various questionable assumptions underlying it are granted. Neither does Adams and Garrison's alternative criterion. In fact, extended cognitive systems easily satisfy these criteria. Finally, it was shown that Sprevak's objection to the extended mind thesis rests on a sleight of hand. There is simply no good reason *not* to think that minds are, at least sometimes, extended.

<center>**V. Extended Consciousness?**</center>

1. Introduction

In the previous chapter, I defended the extended mind thesis, according to which the physical basis of mental states and processes can and sometimes does extend beyond the boundaries of brain and body by virtue of agents' couplings with external objects. In chapter II, I defended a sensorimotor approach to visual perception, which emphasises the ways in which perceivers' embodiment shapes the content of their experience. Despite some prima facie affinity between the extended mind thesis and the sensorimotor approach, sharing as they do a heavy emphasis on the centrality of action, embodiment, and the contributions of the surrounding environment to certain mental phenomena, the two camps are currently at war over the issue of the localisation of consciousness. In a nutshell, the central question driving the debate is this: is conscious experience something that happens in the world or inside the head?

The examples of extended cognition typically found in the extended mind literature - playing Tetris, computation distributed across pen and paper, notebooks as external memory sources, and the like - tend to focus on low-level, non-conscious aspects of mentality. It is, after all, common wisdom in both contemporary cognitive science and philosophy of mind that a great deal of cognitive processing occurs below the threshold of conscious awareness. Because of this, there are many defenders of the extended mind thesis who remain keen to draw a sharp line between these sorts of mental phenomena, which they are happy to admit may be constituted by physical processes not confined solely to the brain and full-blown conscious experience which they insist, for reasons explored below, is still solely the accomplishment of neural

<center>207</center>

processes locked away inside the subject's skull. Even for some philosophers impressed by the arguments in favour of extended cognition, the notion of extended consciousness is just one step too far. According to what I shall call the *extended conscious mind thesis*, however, the physical basis of consciousness is not confined exclusively to the brain either, but incorporates parts of the world in which it is situated also. This is the topic of this chapter.

The extended conscious mind thesis is enjoying growing support among philosophers inspired by both the sensorimotor approach to perception and the extended mind thesis. Although it has obvious parallels with the extended mind thesis, the extended mind thesis' originator and most prominent defender (as well as the hero of my previous chapter), Andy Clark, argues that extended mind theorists can and should reject the extended conscious mind thesis, and offers an alternative internalist account which admits a causal but non-constitutive role for bodily engagements with the world. In this final substantive chapter I examine how well this claim fits with some of the key commitments of his extended mind thesis, and the implications for the extended mind theorist who wishes to deny the extended conscious mind thesis. I argue that Clark's position is unsustainable; his extended mind theorist does not, in fact, have the resources to reject the extended conscious mind thesis, and the extended conscious mind thesis deserves to be taken far more seriously.

2.   The extended mind thesis and the extended conscious mind thesis

As we have seen, the sensorimotor approaches to perception characterises perceptual experience as a temporally extended bodily engagement with the environment,

mediated by practical understanding of the effects of movement on sensory experience. One alleged philosophical consequence of this approach (by, e.g. Ward 2012 and Noë 2004 ch. VII) is that the 'vehicles' that physically realise conscious experience might not be confined to the brain. Some authors urge that if it is indeed the case that perceptual experience is closely dependent on bodily engagements with the world, then the implicated patterns of environmental interaction should count among the physical constituents of perceptual experience. Following a convention found in the sensorimotor theory literature, I shall refer to these patterns of environmental interaction as *sensorimotor dynamics*.

One of the most ardent defences of the extended conscious mind thesis can be found in the closing chapter of Noë's *Action in Perception:*

> What I have been defending […] is externalism about the vehicles of content of experience. I have been arguing that, for at least some experiences, the physical substrate of the experience may cross boundaries, implicating neural, bodily and environmental features. (Noë 2004 p.221, italics removed)

> [B]rain, body and world work together to make consciousness happen […] Experience is not caused by and realised in the brain, although it depends causally on the brain. Experience is realised in the active life of the skilful animal. (ibid. p.227)

And something very much like it is heavily implied, if not explicitly endorsed, by Noë's former collaborator, O'Regan, in his latest book[55]. For example, O'Regan writes:

---

[55]An *Adaptive Behaviour* referee suggested that attributing the extended conscious mind thesis to O'Regan slightly misrepresents his position, as his developed view is

Admittedly there is some information being built up in the brain, namely the information about all the things I can do with my eyes and the changes that will occur as I move them around. You could argue that this information constitutes a kind of representation. But the *seeing* is not caused by the existence or "activation" of this information or representation. It is caused by –I had better say *constituted by*- being engaged in making use of this information to explore or *visually manipulate* the visual scene. (O'Regan 2011. p.28, italics in original)

And later:

If we abandon the idea that feels [O'Regan's term for the qualitative aspects of sensations] are the kind of thing that are generated, and instead take the view that they are constituted by skilled modes of interaction with the environment, then the […] mysteries about feel dissipate. Indeed we can explain why feels have these mysterious characteristics. (ibid. p.115, my italics)

Versions of the extended conscious mind thesis are also defended by (inter alios) Kiverstein & Farina (forthcoming), Loughlin (2012), Ward (2012), Manzotti (2011), Rowlands (2010a; 2003), McCulloch (2003), and Hurley (1998), and are seriously

---

that perceptual consciousness consists in higher order cognitive access to what he calls 'feels' - the qualitative aspect of experience. I disagree. Although O'Regan maintains that higher order cognitive access is necessary for a subject to be conscious of a feel, and this is presumably the work of certain neural mechanisms, this doesn't change the fact that on his view, the 'what-it-is-like-ness' of vision and other senses is not internal to the brain, but is constituted by sensorimotor interactions, as he makes apparent with his analogy between visual qualities and the 'spongeyness' of a sponge. (O'Regan 2011 ch. VIIII).. Hence, there is an undeniable element of externalism inherent in his account.

considered, though not fully endorsed, by Hutto and Myin (2013), Wheeler (2010), Chalmers (2008) and Wilson (2004).

Note that the extended conscious mind thesis does *not* claim that the vehicles of all forms of consciousness extend into the environment. Its proponents can (and do: e.g. Noë 2004 pp. 213-4) readily accept the existence of forms of consciousness the material basis of which are entirely internal/neural. The extended conscious mind thesis pertains to *perceptual experience* in particular. In keeping with recent literature, I shall focus exclusively on vision, though I see no principled reason why similar morals might not be drawn, mutatis mutandis, for at least some other sensory modalities, such as touch. Head and Holmes' (1911) classic example of the blind man using a cane to directly perceive the path in front of him, discussed previously, could arguably provide one such example. I will therefore be using 'perception' as shorthand for 'visual perception' throughout. Other forms of extended consciousness may be possible, perhaps even actual, but for reasons of brevity these are not examined here. To be absolutely clear, note also that the extended conscious mind thesis certainly does not claim that the brain is not necessary for conscious perception, nor does it deny that the brain is the most important or most scientifically challenging part of consciousness' physical substrate, only that it is not its sole physical constituent. Just as, contra Adams and Aizawa, no extended mind theorist would claim that a pencil coupled to a mathematician thinks about sums, no extended conscious mind theorist would claim that an object in a perceiver's surroundings is itself a subject of experience.

Recall that according to the extended mind thesis, cognition can and frequently does extend beyond the brain to incorporate structures in a subject's environment. Extended mind theorists are wont to claim that external objects like notebooks or

mobile phones can partly constitute the physical basis of a mental state such as a belief, providing they are suitably coupled to an agent. For instance, the information therein must be readily available, generally reliable, automatically endorsed upon retrieval, and guide the agent's behaviour in the relevant way (Clark & Chalmers 1998 p.17). Similarly, the extended mind thesis has it that subjects' capacities for action and perception allow them to use information in their surroundings in place of detailed internal representations by offloading some of the brain's workload onto the environment, thereby expanding the physical basis of cognitive processing. Clark and Chalmers illustrate this claim using the well-known experiments by Kirsh and Maglio (1994) showing that players of the retro videogame Tetris, rather than acting on prior decisions based on the formation and rotation of mental images, identify the best shape of fit for falling puzzle pieces on a computer screen by manipulating the shapes themselves. On Clark and Chalmers' interpretation of Kirsh and Maglio's data, the physical basis of the thought process by which the puzzle is solved literally comprises the shapes displayed on the screen and the actions by which they are manipulated.

It might appear reasonable to suppose that if the extended mind thesis is accepted then similar morals could apply with respect to sensorimotor dynamics and conscious perception. For example, Wilson (2010) points out that Clark and Chalmers's own criteria for something's being coupled to an agent, and therefore comprising an extended cognitive system, are satisfied by perceivers' environments. We have a tendency to automatically endorse our basic perceptual beliefs (ceteris paribus). These guide our action, and are reliably caused by our environment, which is readily available for visual inspection. And as Menary (2010c; 2007) stresses, the vast majority of cognitive couplings arise from sensorimotor engagements, suggesting a

strong prima facie affinity between the sensorimotor approach of perception and the extended mind thesis. Such an affinity is clearly evidenced in the empirical literature on vision by O'Regan's (1992) characterisation of the external world as an external memory store accessed by visual attention. Clark, however, resists the application of extended mind-style reasoning to the topic of consciousness. Though Chalmers has since expressed some openness, at least in principle, to the idea of extended consciousness (Chalmers 2008), in their original *Analysis* paper he and Clark stress that their arguments are intended to apply only to unconscious, low-level cognitive processing only

> Some find this sort of externalism unpalatable. One reason may be that many identify the cognitive with the conscious, and it seems far from plausible that consciousness extends outside the head in these cases. But not every cognitive process, at least on standard usage, is a conscious process. It is widely accepted that all sorts of processes beyond the borders of consciousness play a crucial role in cognitive processing [...]. So the mere fact that external processes are external where consciousness is internal is no reason to deny that those processes are cognitive. (Clark and Chalmers 1998 p.10)

In response to the burgeoning literature in favour of the extended conscious mind thesis, Clark has recently published two papers arguing that not only do arguments in favour of the extended mind thesis fail to generalise to support the extended conscious mind thesis, but extended mind theorists can and should reject the extended conscious mind thesis in favour of an internalistic alternative whereby sensorimotor dynamics causally contribute to, but in no way constitute, perceptual experience (Clark 2012b; 2009). In what follows I respond to Clark's critique and

defend an interpretation of sensorimotor dynamics in line with the extended conscious mind thesis.

Responses to Clark's original 2009 critique of the extended conscious mind thesis have tended toward one of two methodologies. One tactic has been to appeal to empirical findings that might be taken to support it. Loughlin (2013), for example appeals to neuroscientific research on perceptual binding (specifically Revonsuo (1999)) to argue that there are experimental data on vision which do not straightforwardly lend themselves to internalistic interpretation. And Kiverstein and Farina (forthcoming) offer an account of sensory substitution technologies in line with the extended conscious mind thesis. I am sympathetic to Kiverstain and Farina's conclusion but, as noted in chapter III, I suspect that their conflation of the user's body schema with a body image might prevent their argument from establishing externalism about consciousness. The other tactic involves making a case for a personal-level conception of perceptual experience as active, embodied, and environmentally embedded, then inferring that something like the extended conscious mind thesis follows regarding the sub-personal level. Ward (2012) makes this latter move. However, neither of these addresses the crucial issue of whether or not the extended mind thesis and the extended conscious mind thesis really are as easily divorceable as Clark claims, or engages directly with the framework for understanding the sub-personal machinery of perceptual consciousness Clark has since developed at length (Clark 2013; 2012c) and which he takes to establish his internalism, so in what follows I pursue a slightly different strategy. I argue that given some of the wider commitments of the extended mind thesis and the concessions Clark is prepared to make concerning the contribution of sensorimotor

dynamics to perceptual experience in his 2013 and 2012 papers, a denial of the extended mind thesis is ultimately untenable for his extended mind theorist.

There is one technical difference between the extend mind and extended conscious mind theses worth clarifying before proceeding further. According to certain formulations of the extended mind thesis, including one found in Clark and Chalmers's seminal *Analysis* paper, the mind extends when external media are utilised in a manner coarsely functionally isomorphic to an ordinarily brain-bound cognitive process. This version of the extended mind thesis still allows for – in fact it presupposes – a stable package of internal cognitive capacities which may be extended outwards through intelligent behaviour. From the perspective of the extended conscious mind thesis, there is no analogous stable package of visual consciousness to extend *from*. Rather, because perception consists primarily in a dynamic sensorimotor engagement with the world, it extends beyond the brain by its very nature. It is in *this* sense that consciousness can be said to be 'extended'. Again, proponents of the extended conscious mind thesis need not necessarily deny there can be perception-like states which are constituted entirely neurally (hallucinations may be one such example). The point is that in the active exploration of an environment, the physical process underling perceptual experience is not just *caused* by dynamic interactions between agent and environment, it *includes* them.

3.   **The horizonal structure of perception and action-oriented predictive processing**

As aforementioned, the sensorimotor approach to visual perception tends to get started with the phenomenological fact that perceptual experience outstrips what is

strictly present to the eyes at any given instant. Despite not being located in the visual field, occluded sides and unattended perceptible features of objects are experienced as present and accessible – we anticipate further presentations of that object (for an insightful early discussion of this phenomenon, see Husserl 1999b). I have called this anticipatory aspect of perceptual experience its *horizonal structure*. The sensorimotor approach cashes out the horizonal structure of perception in terms of perceivers' implicit understanding of the potential changes in experience that would result from further exploration of their environment through action. On such a view, the occluded objects are understood implicitly in experience as 'available to perception through appropriate movement' (Noë 2012 p.58, italics removed). Experience is therefore construed not as a linear succession of static representational states, but as a temporally extended process of environmental probing. On an interpretation in line with the extended conscious mind thesis, the exercise of such embodied capacities and the implicated sensorimotor dynamics are as constitutive of the experience as the accompanying brain processes. Clark resists this conclusion, urging instead that sensorimotor dynamics are not constitutive of experience, but merely causally impinge upon what is genuinely constitutive of experience, namely certain structures in the brain. Hence, despite being a renowned champion of externalism about the vehicles of cognition, Clark is a conservative internalist about the vehicles of consciousness

> I do think that there is something right about [the extended conscious mind thesis]. What is right is the idea that experience, as it unfolds in most normal daily (awake) circumstances, is directly world-revealing, and involves a crucial and complex dance between sensory transduction and real-world action […] But it does not follow that the material (sub-personal processing)

basis of that experience must then extend to encompass its own objects.

Instead, the material apparatus can still quite reasonably be thought to be

wholly internal, consisting in the way neural systems both elicit and respond

to signature perturbations in the environment. (Clark 2012b p.3)

This position is fleshed out using the action-oriented predictive processing

framework Clark outlines in detail elsewhere (2013; 2012c). Drawing from an

impressive and convincing range of empirical studies using predictive coding[56] to

study the visual cortex (e.g. Rao & Ballard 1999), Clark suggests that the horizonal

structure of perceptual experience is best understood in terms of the brain's capacity

to predict its own sensory states using stored knowledge (learned or innate, or a

combination of the two) of generative models of the external world. Incoming

sensory signals are modulated by sensorimotor dynamics and the top-down

predictive processing is refined through learning in a familiar connectionist style by

means of the back propagation of error[57]. Clark explains, 'The key idea is that the

brain uses prediction-driven processing routines to acquire and deploy hierarchical

generative models of the hidden causes (sometimes called latent variables) that best

explain the changing patterns of sensory input that impinge upon the agent' (Clark

2012b. p.8). By acting in such a way as to drive along the relevant neural process by,

for example, turning one's head or eyes or an object closer to one's body, the

perceiver plays an active role in structuring the flow of information responsible for

the content of their (internally realised) perceptual experience, which incorporates an

expectation of further presentations of whatever the subject is currently looking at.

Borrowing a term from Pfeifer et al (2007), Clark characterises this intelligent use of

---

[56] Software which makes uses a relatively small amount of inputted information, to make accurate generalisations or 'predictions' about a wider range of information.

[57] I.e. Using current input to retroactively correct errors.

action to induce one's own sensory content as a form of *informational self-structuring* (Clark 2012b p.10). The intended upshot of the application of the action-oriented predictive processing framework is to accommodate the contributions of sensorimotor dynamics to perceptual experience, and account for perceptual experience's horizonal structure, without conceding the truth of the extended conscious mind thesis. As Clark puts it, 'Embodiment and action fit very naturally within such a framework. Embodiment matters because embodied agents are active agents, and active agents can systematically alter their own sensory input streams in ways that can drive faster and more successful prediction-based learning' (ibid).

For the remainder of this chapter, I grant that the horizonal structure of perceptual experience owes to the brain's capacity to predict its own sensory states. Indeed, not only does Clark (2012c) present an impressive wealth of compelling empirical evidence for this, its truth would be conducive to the extended conscious mind thesis and the sensorimotor approach generally, as it puts the horizonal structure of perception on a firm neurocomputational footing hitherto absent from the discussion. Having defended it at great length in the previous chapter, I also grant the truth of the extended mind thesis, that is, that under certain conditions cognitive processing occurs partly beyond the boundaries of brain and body. What I shall question, and ultimately reject, is the claim that the extended mind theorist who endorses the action-oriented predictive processing framework can coherently and unproblematically reject the extended conscious mind thesis.

4. Perceptual cognition and informational self-structuring

As the extended conscious mind thesis concerns perception, one obvious tactic for refuting it would be to argue that unlike the extended cognitive processing described by extended mind theorists, the cognitive processing underlying perception does not extend. This would nip the extended conscious mind thesis in the bud. As discussed previously, in earlier writings Clark has explicitly denied cognitive extension in the case of perception by appealing to Milner and Goodale's dual streams hypothesis, according to which vision and action have separate underlying cognitive architectures, to drive some conceptual elbow room between vision and sensorimotor dynamics. These arguments notwithstanding, Clark's 2012b paper represents something of an unacknowledged concession. It is an unnoticed consequence of Clark's position that the extended mind theorist who endorses the action-oriented predictive processing framework cannot deny extended perceptual cognition on pain of inconsistency. This is because the centrality of informational self-structuring to the action-oriented predictive processing, when considered in light of some key commitments of the extended mind thesis expressed by Clark elsewhere, unwittingly commits the extended mind theorist to an extended physical basis of perceptual cognition.

In his book *Supersizing the Mind*, Clark presents informational self-structuring as a paradigm case of cognitive extension in action, as a central tenet of his version of the extended mind thesis is that actions, when undertaken to alleviate some of the brain's workload for the sake of improved cognitive performance, qualify as partly constitutive of that cognitive processes. Clark appeals to experimental work by the likes of McNeil (2005) and Goldin-Meadow (2003) on gesture to support an extended mind-style interpretation of informational self-structuring (Clark 2008 pp.123-35. cf. Wheeler 2013). These studies teach us that gestures both encode

information over and above what is represented in the brain, as occurs in cases of 'speech-gesture mismatch' (see Goldin-Meadow & Singer 2003), and play a crucial role in the information transformation required for self-expression and consequently communication. Clark repeatedly asserts that only dogmatic internalistic prejudices disbar us from counting gestures as partly constitutive of the cognitive processing in which they are implicated

> [...] the key distinction between "merely impacting" some inner cognitive process and forming a proper part of an *extended* cognitive process looks much less clear (...) in cases involving the systematic effects of *self-generated* external structure on thought and reason. (Clark 2008, p.126 italics in original)

> These are cases when we confront a recognisably cognitive process, running in some agent, that creates outputs (speech, gesture, expressive movements, written words) that, recycled as inputs, drive the cognitive process along. In such cases, any intuitive ban on counting inputs as parts of mechanism seems wrong. Instead, we confront something rather like the cognitive equivalent of a forced induction system. (ibid. p.131)

As these quotations clearly indicate, informational self-structuring is not confined to gesture. To give but one example, complementary findings come from Villegas, Castro and Gutierrez (2009), who found that undergraduate mathematics students' verbalisations significantly increased the efficiency with which they solved a set of optimisation problems[58]. The students quite literally boosted their own thought processes by talking to themselves about the steps they were taking to solve the

---

[58] The task of finding the best feasible solution to a given mathematical problem.

problems. In light of Clark's take on informational self-structuring, it follows that for the extended mind theorist these and analogous performances, by functioning to self-structure internal processing, also qualify as genuine constituents of cognitive processing. On the anodyne assumption that perceptual processing is part and parcel of cognitive processing, the same must go for the environmental interactions required by the action-oriented predictive processing framework, for these are also deliberate intelligent performances undertaken for the sake of driving along one's internal cognitive processing and alleviating the need for the detailed bottom-up representations of the external world envisioned by the likes of Marr (1982) (Clark 2013 p.2). Hence, contra Clark, the extended mind theorist *is* committed to sensorimotor dynamics being partly constitutive of perceptual cognition.[59] This is, in effect, precisely the sort of conception of perceptual cognition advocated by sensorimotor theorists like O'Regan

> Seeing and perceiving are not achievements of an isolated head or brain, quietly humming along on its own. The organism moves its eyes, repositions its body to get a better perceptual grip on the objects that surround it, and thereby attempts to advance in the execution of the hierarchy of ongoing projects it is engaged in. *The locus of perceptual processing includes the world rather than being just confined to the head*. (O'Regan & Myin 2009, italics added)

To reiterate: according to the action-oriented predictive processing framework, the horizonal structure of perception is best explained by the brain's apparent ability to anticipate its own sensory states, and this process is driven along by perceivers'

---

[59] Not all extended mind theorists are hostile to the idea of extended perceptual cognition. Rowlands, a defender of both the extended mind and extended conscious mind theses, argues for such a view (2010a; 2010b pp.107-34).

bodily interactions with the world by means of actions such as head-turnings, eye saccades, and generally manoeuvring oneself in relation to the object. As such, perceivers actively induce the content of their own perceptual experience by self-structuring their internal cognitive-sensory states. Following Clark's own extended mind characterisation of informational self-structuring, then, these interactions with the objects of perception are also constitutive parts of the mechanisms of perceptual cognition. The extended mind theorist who endorses the action-oriented predictive processing framework is therefore barred from denying extended perceptual cognition in order to reject the extended conscious mind thesis outright. If extended perceptual cognition is granted, what is required of the extended mind theorist who still wishes to reject the extended conscious mind thesis is a principled reason for distinguishing between the parts of these extended perceptual mechanisms constitutive of conscious experience, and what we might call the 'merely cognitive' parts which are not, without which any such distinction will be arbitrary and possibly question-begging. The remainder of the paper considers the extent to which the extended mind theorist who wishes to maintain internalism about consciousness can meet this challenge.

5. The coupling-constitution fallacy revisited

In his initial treatment of the extended conscious mind thesis prior to publishing his formulation of the action-oriented predictive processing framework, Clark (2009) accuses its defenders of deploying the sort of erroneous reasoning attacked by Adams and Aizawa's coupling-constitution argument. Recall that the fallacy they attack, the 'coupling-constitution fallacy', is

[A] tacit move from the observation that process X is in some way causally connected (coupled) to a cognitive process Y to the conclusion that X is part of the cognitive process Y. The pattern of reasoning here involves moving from the observation that process X is in some way causally connected (coupled) to a process Y of type ϕ to the conclusion that X is part of a process of type ϕ. (Adams & Aizawa 2009 p.81)

In the context of the debate over extended consciousness, the alleged fallacy consists in inferring from the fact that sensorimotor dynamics impact upon conscious experience that they are therefore constitutive of that experience (Clark 2009 p.981). Indeed, there are examples of defenders of the extended conscious mind thesis apparently committing this fallacy. Here is a particularly unfortunate passage from Noë: 'if ever there was a plausible candidate for a psychological state that is *driven and so partially constituted by the environment,* it is perceptual consciousness' (Noë 2008 p.460, quoted in Clark 2009 p.982, italics in Original). Elsewhere, as noted in the previous chapter, Noë even attributes such a view to Clark and Chalmers in what appears to be a severe misreading of their case for the extended mind thesis: 'according to [the extended mind thesis], the environment can *drive and so partially constitute* cognitive processes' (Noë 2004 p.221, my italics). Here Noë misrepresents Clark and Chalmers' case for extended cognition, which appeals to functional isomorphism between internal and external information-carrying structures, rather than any inference from coupling to constitution, then recasts it as an argument for extended consciousness.

Clark is certainly correct to point out that as with the extended mind thesis extended consciousness cannot be established purely on the basis of an inference from causal interaction to constitution. But trivially, however the mind works and wherever and

however it is physically realised, it will encompass a range of causally interacting

physical states and processes. So there is seemingly no a priori obstacle to the

defender of the extended conscious mind thesis simply chalking up these instances of

the coupling-constitution fallacy to careless wording on Noë's part and recasting

their thesis purely in terms of constitution. And as Clark himself notes, one major

problem with debates over the boundaries of the mind is that opposing arguments

often depend heavily for their plausibility on extant internalist or externalist

prejudices, resulting in dialectical stalemate (Clark 2009 p. 981). In an attempt to

block such a move and alleviate this threat of stalemate, Clark offers an empirical

argument for demarcating the constitutive from the casually related in terms of the

internal (neural) and external (bodily/worldly) for *consciousness* in a way that differs

from the standard coupling-constitution arguments against EM.

We have seen that coupling-constitution arguments against the extended mind thesis

presuppose a prior account of cognitive demarcation, without which the distinction

between what does and does not count as constitutive of cognition within a coupled

system is arbitrary. Hence the controversy over the 'mark of the cognitive'- an

elusive criterion that would ground the causal/constitutive distinction by providing a

necessary and sufficient condition (or conditions) for a state or process to qualify as

genuinely *cognitive*. By parity of reasoning, then, an analogous 'mark of the

conscious' is required to ground a coupling-constitution argument against the

extended conscious mind thesis. Clark's candidate for a mark of the conscious is

processing power. Adapting an argument due to Eliasmith (2009), Clark suggests

that unlike the brain with its billions of neurons and trillions of synapses, non-neural

states and processes have nowhere near the bandwidth required to support a rich

stream of conscious experience. For example, nothing at all like the 40-70Hz

oscillations in the cerebral cortex purportedly demonstrated as necessary for perceptual binding by Crick and Koch (1990) occurs outside of the brain. This is supposed to give reason for thinking that it is a contingent empirical fact that consciousness, unlike low-level nonconscious cognitive processing which need not be subject to the same bandwidth constraints, is constituted entirely neurally

> [I]n the special case of arguments for [the extended conscious mind thesis], we can begin to discharge this obligation [to provide a criterion of demarcation]. For if indeed the physical machinery of consciousness requires fast timescale operations and processing, and the non-neural body acts as a low-pass filter preventing external (…) signals from directly entering into such operations and processing, then such signals are fit to play only a causal role, driving the neural systems within which the right kinds of fast binding and processing can occur. In such cases, one might all manner of complex couplings without thereby producing an extended material base for conscious experience. (Clark 2009 p.986)

It seems uncontroversial that nothing like fast high-bandwidth neural processing occurs on the non-neural side of perceptual processing, but there is reason to be sceptical as to whether this fact alone suffices to ground the causal-constitutive distinction the extended mind theorist who wishes to deny extended consciousness needs to neatly distinguish the conscious from the 'merely cognitive' parts of the extended perceptual mechanisms they are forced to concede. While certain neural features will doubtless be *necessary* for conscious experience to occur, we cannot infer from this fact alone that they are *sufficient* to secure the phenomenology of all or even any form of consciousness (a point made by Ward 2012). At best, this only shows that the relevant fast neural processes must be implicated in every mechanism

underlying an experience; it does not prove that these neural processes are the sole physical constituent of that experience, nor does it automatically licence the inference that all other parts of a mechanism in which a state that has that feature is implicated are ipso facto non-constitutive of the associated experience.

The fact that high bandwidth neural processing is required for perceptual experience does not by itself establish internalism about consciousness. Merely pointing to the fact that consciousness requires fast neural processing does not suffice to provide a mark of the conscious. What is required is an account of what neural processes actually *do* that renders the wider mechanisms in which they are embedded constitutively redundant. Clark's second paper against the extended conscious mind thesis (Clark 2012b) invokes details of his action-oriented predictive processing framework, absent from his 2009 treatment, which he takes to do just this. Hence, in the next section I consider whether or not appealing to the action-oriented predictive processing framework can establish internalism, answering in the negative.

6.   The stuff dreams are made of?

Some predictive coding enthusiasts take the results of its application to the scientific study of vision as evidence for an indirect theory of perception (e.g. Hohwy 2007). Perception is said to be 'indirect' if its objects are not the actual things in the world experience purports to present us with, but mental intermediaries like 'sense-data' (see e.g. Moore 1953, Russell 2001) or something similar (e.g. Jackson 2004), and 'direct' if its objects are not mental intermediaries, but the objects themselves (see, e.g. Stoneham 2008, Martin 2002 and Austin 1962). In this instance, the intermediary mental objet would be 'the brain's best hypothesis, as embodied in a high-level

generative model, about the causes in the outer world' (Hohwy 2007 p.323).

Accepting the indirectness of perception as a consequence of the action-oriented

predictive processing framework would be a quick and easy way of establishing

internalism about perceptual consciousness. If the world of experience comprises

mental objects locked away inside the skull, then it makes little if any sense to

suggest that the physical basis of that experience extends beyond it. But Clark is

eager to distance himself from such a view, insisting that perception according to the

action-oriented predictive processing framework would be better described as 'not-

indirect perception' (2013 p.27), whereby stored generative models are not the

objects of perception, but sub-personal states which enable us to perceive the world

itself

> [W]e may still reject the bald claim that "what we perceive is the brain's best
>
> hypothesis," Even if our own prediction is indeed (…) doing much of the
>
> heavy lifting, it remains correct to say that *what* we perceive is not some
>
> internal representation or hypothesis but (precisely) the world. We do so
>
> courtesy of the brain's ability to latch on to how the world is by means of a
>
> complex flow of sub-personal processes. (Clark 2012c p. 54)

Having dismissed as erroneous and 'Cartesian' any dichotomy between the world of

experience inside the head and the actual world outside of it, much like sensorimotor

theorists, what is left to motivate Clark's internalism? For Clark, it is the 'super-tight

empirical link' (Clark 2012b p.18) between perception and phenomenologically

similar yet paradigmatically internally-constituted experiences like imagination and

dreaming. The quasi-visual imagery of dreams and other imaginings exhibits a

horizonal structure much like that of perception, which the action-oriented predictive

processing framework explains in terms of a common underlying cognitive

architecture. The basic idea is that in waking perception, when all goes well, the brain matches the appropriate generative model to its distal cause(s) by adjusting for error in the incoming sensory signal (Clark 2012c pp.44-55). In forms of vision-like experience which fall short of perception, there is no stable incoming sensory signal to fine-tune the generative model. So despite their shared architecture, perceiving and dreaming can differ phenomenologically

> [N]othing in [the action-oriented predictive processing framework] requires that the system, when simply cycling, in the sleep or imagining state, in the absence of ongoing driving external inputs, will typically support the very same kinds of stability and richness of experienced detail that daily sensory engagements offer. (Clark 2012b p.16)

These differences in stability and richness are worth elaborating upon. Contrary to the well-worn fantasies of epistemological scepticism and popular science fiction, it is in fact fairly easy to tell whether or not one is dreaming, though doing so can require a period of training. The phenomenon of lucid dreaming testifies to this. Though lucid dreaming can be made to sound like pseudo-spiritualist mumbo-jumbo, and is often unfortunately presented as such in popular culture, there is in fact nothing anti-physicalist or mystical about it, and it is a subject of serious scientific study by psychologists, neuroscientists and medical practitioners (for a thorough scientific review, see Hobson 2009). Though experimental studies of lucid dreaming are unavoidably plagued by the sometimes-unreliable methodology of first-person report, there is widespread agreement that the phenomenon is genuine. Lucid dreaming occurs when people become aware that they are dreaming and regain a sense of agency within their dream. Lucid dreams usually happen (that is, when they do) during REM sleep (LaBerge 1990), and there is evidence that they tend to occur

more frequently in children than adults (Voss et al 2012). However, adults can become lucid dreamers by learning to recognise telltale signs that they are dreaming by routinely performing 'reality checks' while awake. These involve deliberately attending to some features of the world that are largely absent from dreams. For example, the constancy of perceived properties of objects and the stability in the layout of one's environment present in perception break down in dreams. Lucid dreamers learn to distinguish dreams from waking perception on this basis by habituating themselves to the practise of regularly attending to features of their surroundings like the contents of containers, the reliable compliance of contraptions like windows and doors, and the consistent readability of texts, signs and clock faces. Following a period of regular interrogation of the surrounding world several times a day, these reality-checking routines are instinctively initiated when 'cycling' in sleep, allowing skilled reality-checkers to notice these phenomenological discrepancies as they unfold (Metzinger 2009 pp.133-48, LaBerge 2004)[60]. With enough practice, lucid dreamers can even communicate with experimenters while dreaming through a previously agreed upon series of eye movements prompted by conditioned stimuli, usually light (LaBerge 1990).

Lucid dreaming research teaches us that dreaming is not just phenomenologically different to perceiving, it is phenomenologically *impoverished*, or rather: dreaming differs phenomenologically from perceiving *because* it is relatively impoverished.

---

[60] In the spirit of 'experimental philosophy', I have experimented with these techniques myself with some success, though I often find myself waking up shortly after the moment of realisation in a kind of 'short-circuiting' effect. My favourite finding thus far has been that the removal of my glasses in a dream (the presence of which is interesting in itself, seeing as I've only worn them for a few years) made no difference to the quality of the image.

This phenomenological insight is discussed at length by Jean-Paul Sartre in *The Imaginary,* where he argues that mental images are characterised by an 'essential poverty' (Sartre 2004 p.16) that is not shared by the ordinary objects of perception. Much of Sartre's reasoning on this subject has been rehearsed more recently in the contemporary analytic tradition by Colin McGinn (McGinn 2004). Similarly, O'Regan and Noë remark

> A hallmark of dream-like experiences is the unstable and seemingly random character of dreamt detail...This suggest that without the world to serve as its own external model, the visual system lacks resources to hold an experienced world steady. (O'Regan and Noë 2001 p.947)

The key observation here is that dreams lack the rich detail and reliable structure proper to waking perception, and the fact that it can take a modicum of training to appreciate this is no more a counterargument than somebody's inability to distinguish a Metallica song from an Iron Maiden song is a counterargument to the (perfectly truthful) assertion that they sound completely unalike.[61] Note that, as with differentiating two pieces of music, differentiating dreaming from perceiving cannot always be done instantaneously or over an arbitrarily short period of time, or indeed without the required skill. If all one is allowed to hear is an E5 chord and a single hit of a snare drum, then the differences between the two songs may be present and available but still go unnoticed (note that there may still be detectable differences here - in guitar tone or production style, for example, - to the suitably skilled listener). Similarly, the differences between the experiences of glancing at an imaginary table for two seconds in a dream and glancing at an actual table for two seconds can go unappreciated. But, if one knows what one is looking for, the

---

[61] This example is adapted from Noë's discussion of hallucination (Noë 2004 p.80).

phenomenological differences will be made salient through a slightly longer process of exploration and interrogation. If perception is 'not-indirect', then these phenomenological differences between dreaming and perceiving are explained by what is in the world. Again, O'Regan's (1992) characterisation of the perceived world as an external working memory is instructive in this regard, as the stability of waking perception is explained not in terms of sameness of intermediating representational content, but in terms of the stability and continual availability of information in the world itself.

Returning to the debate over the extended conscious mind thesis, at this juncture the threat of dialectical stalemate resurfaces. Both advocates of the extended conscious mind thesis and extended mind theorists who reject it can agree that the phenomenological differences between perception and dreaming owe to the presence (in waking perception) or absence (during sleep) of the extended sensorimotor dynamics, modulated by the active informational self-structuring needed to fine-tune the brain's generative model. In other words, the poverty of the experience owes to a poverty of sub-personal mechanism. With extended perceptual cognition now in place for both extended conscious mind and extended mind theorists, and having ruled out an appeal to the mark of the conscious as a clear cut deciding factor, the onus is on the extended mind theorist who rejects the extended conscious mind thesis to show why these important phenomenological differences in richness and structure are not constituted by extended sensorimotor dynamics *as well as* the tokened generative model.

How are we to decide what does and does not deserve constitutive status with respect to consciousness? Having deliberately bracketed what he dubs 'more metaphysical' versions of this question (Clark 2012b; 2009), Clark's chosen (though largely

implicit) strategy is to identify reliable correlations between conscious experiences and the sub-personal physical states and processes which accompany them. Hence, the job required of the action-oriented predictive processing framework, if it is to establish internalism, is to provide an answer to what Chalmers (1995; 1996) calls an 'easy problem', as opposed to the 'hard problem', of consciousness. That is, rather than addressing the perplexing puzzle of why or how any physical process should give rise to any experience *at all*, the extended mind theorist who wants to hang onto internalism about consciousness needs only to indicate which physical processes are the most suitable candidates for the vehicles of experience. If the presence of a particular physical process secures the phenomenology, then it is extraneous to include wider processes in which that state or process is embedded in a description of the physical basis of the experience in question.

If this is the approach taken, however, the extended conscious mind thesis will win out. To see why, consider a simple internalistic analogy. Imagine a neuroscientist investigating the neural correlates of a particular experience. Here I follow Chalmers' instructive definition of a 'direct' neural correlate of consciousness

> N will be an NCC when (1) the states of N suffice for the corresponding states of consciousness, and (2) no proper part M of N is such that the states of M suffice for the corresponding states of consciousness. In this way, we pare down any potential NCC to its core: any irrelevant material will be whittled away, and an NCC will be required to contain only the core processes that suffice for the conscious state in question [where N is a neural state and NCC is a neural correlate of consciousness]. (Chalmers 2000 p.25)

Having identified a particular pattern of neuronal spiking in a test subject as correlating reliably with the experience in question, the neuroscientist then attempts to artificially recreate that experience by applying trans-cranial magnetic stimulation or some similar technique to the identified cortical region(s). Let's say that the experiment is successful and that by reproducing the previously identified neural state in the test subject, the experience is reproduced in its entirety. Prompted by doubts that she may have oversold the contribution of one part of the brain to the experience under investigation, the neuroscientist then performs a second test, stimulating only an isolated section of the previously identified cortical region(s). When quizzed, her test subject reports that the induced experience differs phenomenologically to the original experience. The reasonable conclusion for the neuroscientist to draw at this point is that as the original experience is not reproduced by stimulation of only part of the initially identified cortical region(s), that part of the cortex, though implicated in the original experience, is not properly speaking its correlate; there is something about the simultaneous presence of – or more likely the interaction between – the isolated region and the rest of the initially identified cortical area(s), that determines the phenomenology of the original experience. The isolated cortical area taken in itself is the correlate of an experience – a different but similar experience – but the neural correlate of the original experience must be the wider pattern of cortical activity identified initially, as this is what ensures the phenomenology.

This is, admittedly, a philosopher's caricature of neuroscientific research, but it's not too much of a distortion as to be uninstructive (see e.g. Tong 2003, and the essays collected in Metzinger 2000 for some actual examples of this sort of approach in action). Assuming it's a reasonably accurate caricature, however, we can draw some

tentative morals regarding extended consciousness and Clark's action-oriented predictive processing framework. By parity of reasoning, for internalism to remain plausible given the apparent necessity of sensorimotor dynamics and informational self-structuring to rich, stable perceptual experience, its phenomenology would have to be secured solely by the activation of the internal parts of the mechanism responsible for deploying the generative model. But as there apparently are phenomenal features which are occurrent only when the stable sensory signal afforded by sensorimotor dynamics is present, this suggests that the internal machinery of consciousness is only a part (albeit the most important part) of the physical basis of conscious perception. What secures the phenomenology of perception is not the tokening of a generative model, or the self-modulation of a driving sensory signal via the active self-structuring of information, but the dynamic interplay between the two. It follows that the physical correlate of experience – and therefore the best candidate for its vehicle – includes the sensorimotor dynamic in which the relevant predictive neural state is embedded. The action-oriented predictive processing framework is not only compatible with the extended conscious mind thesis, it supports it.

7. Conclusion: putting experience back into the world

I have argued that once it is accepted that sensorimotor dynamics are integral to maintaining perceptual experience, and some basic tenets of the extended mind thesis are granted, internalism about perceptual consciousness becomes increasingly difficult to maintain. The coupling-constitution fallacy loses its bite once it becomes apparent that Clark's action-oriented predictive processing framework commits his

extended theorist to an extended physical basis of perceptual cognition. The extended mind thesis does commit her to this, because it construes actions performed in order to boost cognitive performance as constituents of an extended cognitive process, and the action-oriented predictive processing framework grants a significant role to sensorimotor dynamics which perform precisely this function. The appeal to the brain's processing power as a 'mark of the conscious' fails to ground the causal/constitutive distinction required to demarcate within these extended mechanisms the boundary between the conscious and the merely cognitive, because at best it only provides a necessary, but not sufficient, criterion for the physical realisation of perceptual consciousness. And the brain's tokening of a top-down generative model does not secure the phenomenology of waking perception, whereas the tokening of the model *plus* the active self-structuring of sensory information does, suggesting that, contra Clark, the extended conscious mind thesis *does* hold at a sub-personal level, and not just at a personal level of description given a conception of experience as active, embodied ad environmentally embedded.

In the closing chapter of *Supersizing the Mind*, Clark remarks that minds only appear 'extended' relative to an impoverished picture of cognition as something locked away inside the skull (Clark 2008 p.219). Though it would certainly be premature to deny outright the possibility of formulating alternative arguments for rejecting extended consciousness while holding unto the extended mind thesis, nothing Clark has said shows that the same is not true of conscious experience also.

Conclusions

To summarise: chapter I introduced the concept of the body schema - a system of bodily dispositions and capacities which structure our behaviour and experience, and which exhibits its own particular kind of intentionality. This was contrasted with the notion of a body image, which is a mental state that has the subject's own body as its intentional object. I presented a view of the body schema as flexible, capable of incorporating extra-corporeal objects by virtue of the plasticity of the sub-personal neural representations underlying it, while resisting the identification or reduction of the body schema to these neural representations. Chapter II deployed this notion of the body schema in a defence of the sensorimotor approach to visual perception, according to which perceptual content owes to the possession of sensorimotor understanding. I argued that the body schema plays an essential role in the constitution of this sensorimotor understanding. Chapter III combined the role of the body schema in structuring the content of perceptual experience as outlined in chapter II and the flexibility of the body schema as outlined in chapter I to develop a phenomenological interpretation of perceptual experience via tactile-visual sensory substitution technologies. Chapter IV ventured beyond the body to defend the extended mind thesis, according to which thought itself often takes place partly outside of the head, in the world. Arguments against the extended mind thesis were found to be unconvincing, often relying for their feasibility on undefended internalistic prejudices. Finally, in chapter V, I revisited the sensorimotor approach and considered whether something similar might also be the case with visual consciousness. I argued that, if certain assumptions are granted, then the physical basis of perceptual experience also encompasses components eternal to the brain.

Throughout the thesis, I have gradually pushed the boundaries of consciousness and cognition further and further outwards from the head and into the world, via the body, tacitly advocating a view of the mind as a product of a heterogeneous and highly malleable collection of physical states and processes. Doing so did not necessitate anything as radical as the rejection of any received knowledge about the inner workings of the brain, or an overhauling of the methods of the cognitive sciences. On the contrary, a good deal of mainstream empirical work can be invoked in favour of such a view. Furthermore, and perhaps more importantly from a philosophical point of view, such a view of cognition does better justice to the realities of our ordinary lived experience. Phenomenological analysis, such as that undertaken by Husserl, Heidegger and Merleau-Ponty, teaches us that we are embodied beings at home in a shared world of habitual practices and practical significations. Contemporary cognitive science teaches us that our mental capacities are, like everything else amenable to empirical investigation, constituted out of physical process. Ultimately, given these two most fundamental of facts about our minds, the claim that they may not be firmly locked away behind our eyes should hardly seem all that surprising to begin with.

References

Bibliography

- Adams, F. and Aizawa, K., (2010), 'Defending the bounds of cognition' in R. Menary ed. *The Extended Mind*, MA: MIT Press pp.67-80
- - (2009), 'Why the mind is still in the head' in P. Robbins and M. Aydede eds. *The Cambridge Handbook of Situated Cognition,* NY: Cambridge pp.78-95
- - (2007), *The Bounds of Cognition*, NY: Blackwell
- - (2001), 'The Bounds of Cognition', *Philosophical Psychology* 14 pp.43-64
- Adams, F. and Garrison, G. (2013), 'The mark of the cognitive', *Minds and Machines* 23 pp.339-52
- Armel, K. and Ramachandran, V. (2003), 'Projecting sensations to external objects: evidence from skin conductance response', *Proceedings of the Royal Society of London* 270 pp.1499-1506
- Austin, J. L., (1962), *Sense and Sensibilia,* Oxford: Oxford University Press
- Auvray, M. & Myin, E. (2009), 'Perception with compensatory devices: From sensory substitution to sensorimotor extension', *Cognitive Science* 33 pp.1036-1058
- Baber, C. (2003), *Cognition and Tool Use,* London: Taylor & Francis
- Baddeley, A. (1992), 'Working memory', *Science* 255:5044 pp.556-9
- Bach-y-Rita, P. (2002), 'Sensory substitution and qualia' in. A. Noë and E. Thompson eds. *Vision and Mind*, MA: MIT Press pp.497-515
- - (1972), *Brain Mechanisms in Sensory Substitution,* NY: Academic Press
- Bach-y-Rita, P., Carter, C., Saunders, F., White, B. and Scadden, L. (1969) 'Vision substitution by tactile image projection', *Nature* 221 pp.963-4
- Bach-y-Rita, P., Kaczmarek, K., Tyler, M. and Garcia-Lara, J. (1998), 'Form perception with a 49-point electrotactile stimulus array on the tongue: A technical note', *Journal of Rehabilitation Research and Development* 35:4 pp.427-30
- Bach-y-Rita, P., Kercel, S. (2003), 'Sensory substitution and the human-machine interface', *Trends in Cognitive Science* 7:12 pp.541-6
- Bailey L., (1989), 'Skull's darkroom: The camera obscura and subjectivity' in P. Durbin ed. *Philosophy and Technology: Practical, Historical and Other Dimensions,* Dordrecht: Springer pp.63-79
- Baldwin, T. (2007), 'Speaking and spoken speech' in T. Baldwin ed. *Reading Merleau-Ponty,* London: Routledge pp.87-104
- Baron-Cohen, S. (1996), 'Is there a normal phase of synaesthesia in development?', *Psyche* 2:27
- Bauer, G., Gerstenbrand, F. and Rumpl, E. (1979), 'Varieties of the locked-in syndrome', *Journal of Neurology* 221: 2 pp.77-91
- Bishop, M. and Martin, A. (2014) 'Contemporary sensorimotor theory: A brief introduction' in M. Bishop and A. Martin (eds.), *Contemporary Sensorimotor Theory*, Dordrecht: Springer pp.1-22
- Blakemore, S., Brelstaff, G., Nelson, K. and Troscianko, T., (1995) 'Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes', *Perception* 24 pp.1075-81

- Block, N. (2003), 'Spatial perception via tactile sensation', *Trends in Cognitive Sciences* 7:7 pp.285-6

- Botvinick, M. and Cohen, J. (1998), 'Rubber hand 'feel' touch that eyes see', *Nature* 391 p.756

- Brentano, F. (1995) [1874], *Psychology from an Empirical Standpoint,* trans. A. Rancurello, D. Terrell and L. McAlister, NY: Routledge

- Brooks, R. (1991a), 'Intelligence without reason', *Computers and Thought* online manuscript: http://people.csail.mit.edu/brooks/papers/AIM-1293.pdf (last accessed 22/8/14)

- - (1991b), ' Intelligence without representation' *Artificial Intelligence* 47 pp.139–59

- Buchs. G., Maidenbaum, S. and Amedi, A., (2014), 'Obstacle identification and avoidance using the 'eyecane'' *EuroHaptics* 2014: http://brain.huji.ac.il/publications/Buchs_eurohaptics2014.pdf (last accessed 3/9/14)

- Burge, T. (1979), 'Individualism and the mental', *Midwest Studies in Philosophy* 4:1 pp.73-22

- Cardinali, L., Frassinetti, F., Brozzoli, C., Urquizar,C., Roy, A. and Farnè, A. (2009), 'Tool use induces morphological updating of the body schema', *Current Biology* 19:12 pp.478-9

- Carman, T. (1999), 'The Body in Husserl and Merleau-Ponty', *Philosophical Topics* 27:2 pp.205-26

- Cascio, C., Foss-Feig, J., Burnette, C., Heacock, J., and Cosby, A. (2012), 'The rubber hand illusion in children with autism spectrum disorders: delayed influence of combined tactile and visual input on proprioception', *Autism* 16:4 pp.406-19

- Chalmers, D. (2008), 'Foreword' to A. Clark ,*Supersizing the Mind*, Oxford: Oxford University Press pp. ix-xvi

- - (2000), 'What is a neural correlate of consciousness?', in T. Metzinger ed. *Neural Correlates of Consciousness,* MA: MIT Press pp.17-39

- - (1996), *The Conscious Mind,* Oxford: Oxford University Press

- - (1995), 'Facing up to the problem of consciousness', *Journal of Consciousness Studies,* 2:3 pp.200-19

- Chants., Hindriks, F. and Preyer, G. eds. (forthcoming) *From Individual to Collective Intentionality,* Oxford: Oxford University Press

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

- Chisholm, N. and Gillett G. (2005), 'The patient's journey: Living with locked-in syndrome', *British Medical Journal* 331 pp.94-7

- Chomsky, N. (1959), 'A review of B. F. Skinner's verbal behaviour' *Language* 35:1 pp. 26-58

- Churchland, P. (1981), 'Eliminative materialism and the propositional attitudes', *Journal of Philosophy* 78:2 pp.67-9

- Crick, F. & Koch, K. (1990), 'Towards a neurobiological theory of consciousness', *Seminars in the Neurosciences* 2 pp.263-75

- Clark, A. (2013), 'Expecting the world: perception, prediction, and the origins of human knowledge', *The Journal of Philosophy*, forthcoming (manuscript)

- - (2012a), 'Embodied, embedded, and extended cognition' in K. Frankish and W. Ramsey eds. *The Cambridge Handbook of Cognitive Science,* Cambridge: Cambridge University Press pp.275-91
- - (2012b), 'Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience', *Mind,* forthcoming (manuscript)
- - (2012c), 'Whatever next? Predictive brains, situated agents, and the future of cognitive science', *Behavioural and Brain Sciences,* forthcoming (manuscript)

- - (2010a), 'The extended mind, extended' in R. Menary ed. *The Extended Mind,* MA: MIT Press pp.43-66
- - (2010b), 'Coupling, constitution, and the cognitive kind: A reply to Adams and Aizawa' in R. Menary ed. *The Extended Mind* MA: MIT Press pp.81-99
- - (2009), 'Spreading the joy? why the machinery of consciousness is (probably) still in the head', *Mind* 118: 472 pp.963-93
- - (2008), *Supersizing the Mind,* NY: Oxford University Press
- - (2006), 'Vision as dance? Three challenges for sensorimotor contingency theory', *Psyche* 12:1, online http://www.theassc.org/files/assc/2629.pdf (last accessed 12/8/14)

- - (2003), *Natural Born Cyborgs,* NY: Oxford University Press

- - (2002), 'Is seeing all it seems? Action, reason, and the grand illusion', *Journal of Consciousness Studies* 9:5-6 pp181-202

- - (2001), 'Visual experience and motor action: Are the bonds too tight?', *Philosophical Review* 110:4 pp.495-519
- - (1997), *Being There: Putting Brain, Body and World Together Again,* MA: MIT Press
- Clark, A., and Chalmers, D. (1998), 'The extended mind', *Analysis* 58 pp.7-19
- Clark, A. and Torbio, J. (2001), 'Sensorimotor chauvinism? commentary on O'Regan, J.K., and Noë, A. "A sensorimotor approach to vision and visual consciousness"', *Behavioural and Brain Sciences* 24:5 pp.979-80
- Cochrane, T. (2008), 'Expression and extended cognition', *The Journal of Aesthetics and Art Criticism* 66:4 pp.299-340
- Cole, J. (1995), *Pride and a Daily Marathon,* MA: MIT Press
- Cole, J., Gallagher, S. and McNeil, D. (2002), 'Gesture following deafferentation: A phenomenologically informed experimental study', *Phenomenology and the Cognitive Sciences* 1:1 pp.49-67
- Cole, J., and Paillard, J. (1995), 'Living without touch and peripheral information and body position and movement: Studies with deafferented Subjects' in J. Bermudez, A. Marcel and N. Eilan eds. *The Body and the Self,* MA: MIT Press pp.245-66
- Cowley, S. (2007), 'The cognitive dynamics of distributed language', *Language Sciences* 29:5 pp.575-83
- Crane, T. (2009), 'Intentionalism' in B. McLaughlin, A. Beckermann and S. Walter eds. *The Oxford Handbook of Philosophy of Mind,* Oxford: Oxford University Press pp.474-93
- - (1995), *The Mechanical Mind,* London: Penguin
- Dainton, B. (2005), *Stream of Consciousness,* London: Routledge

- Damasio, A. (2000), *The Feeling of What Happens,* London: Vintage
- Davidson, D. (1984), 'First person authority', *Dialectica* 38:2-3 pp.101-12
- -   (1963), 'Actions, reasons, and causes', *The Journal of Philosophy* 60:23 pp.685-700
- Day, S. (1996) 'Synaesthesia and synaesthetic metaphors', *Psyche* 2:32
- Dennett, D. (2007), 'Heterophenomenology reconsidered', *Phenomenology and the Cognitive Sciences* 6:1-2 pp.247-70
- -   (1996), *Kinds of Minds*, NY: Basic Books
- -   (1991), *Consciousness Explained*, London: Penguin
- -   (1971), 'Intentional systems', *The Journal of Philosophy* 68:4 pp.87-106
- -   (1969), *Content and Consciousness,* London: Kegan Paul
- Deroy, O. and Auvray, M. (forthcoming), 'Beyond vision: The vertical integration of sensory substitution devices' to appear in D. Stokes, M. Matthen & S. Biggs eds. *Perception and its Modalities,* Oxford: Oxford University Press
- Descartes, R., (1998) [1637/1647], *Discourse on the Method and Meditations on First Philosophy,* trans. D. Cress, London: Hackett
- -   (1985) [1629] 'Treatise on man' in *The Philosophical Writings of Descartes,* trans. J. Cottingham, R. Stoothoff and D. Murdoch, Cambridge: Cambridge University Press pp.99-108
- de Vignemont, F. (2014), 'Bodily immunity to error' in. S. Prosser and F. Recanati eds. *Immunity To Error Through Misidentification: New Essays,* Cambridge: Cambridge University Press pp.224-46
- -   (2011), 'Embodiment, ownership and disownership', *Consciousness and Cognition* 20:1 pp.82-93
- -   (2010), 'Body schema and body image - pros and cons', *Neuropsychologia* 48 pp.669-80
- D'Oro, G. and Sandis, C. (2013) eds., *Reasons and Causes: Causalism and Anti-Causalism in the Philosophy of Action,* Basingstoke: Palgrave-Macmillan
- Drayson, Z. (2010), 'Extended cognition and the metaphysics of mind', *Cognitive Systems Research* 11:4 pp.367-77
- Dretske, F. (1981), *Knowledge and the Flow of Information,* MA: MIT Press
- Dreyfus, H. (2000), 'A Merleau-Pontian Critique of Husserl's and Searle's Representationalist Accounts of Action', *Proceedings of the Aristotelian Society* 100 pp.287-302
- -   (1992) *What Computers Still Can't Do,* MA: MIT Press
- Eliasmith, C. (2009), 'Dynamics, control and cognition' in P. Robbins and M. Aydede eds. *The Cambridge Handbook of Situated Cognition,* Cambridge: Cambridge University Press pp.134-54
- Evans, G. (1982), *The Varieties of Reference,* Oxford: Oxford University Press
- Fancher, R. and Rutherford, A. (2012), *Pioneers of Psychology: A History,* NY: W.W. Norton
- Farmer, H., Tajadura-Jiménez, A., and Tsakiris, M. (2012), 'Beyond the colour of my skin: How skin colour effects the sense of body-ownership', *Consciousness and Cognition* 21 pp.1242-56
- Fodor, J. (2008), 'Against darwinism', *Mind and Language* 23:1 pp.1-24
- -   (1990), *A Theory of Content and Other Essays,* MA: MIT Press
- -   (1980), 'Methodological solipsism as a research strategy in cognitive psychology', *Behavioural and Brain Sciences* 3 pp.68-109

- Fodor, J. and Pylyshn, X. (1981), 'How direct is visual perception? Some reflections on Gibson's 'ecological approach'', *Cognition* 9 pp.139-96
- Franz., V., Gegenfurthner, K., and Scharnowski, F. (2005), 'Illusion effects on grasping are temporally constant not dynamic', *Journal of Experimental Psychology* 31:6 pp.1359-78
- Franz, E. and Ramachandran, V. (1998), 'Bimanual coupling in amputees with phantom limbs', *Nature Neuroscience* 1:6 pp.443-4
- Franz, V. (2003), 'Planning versus online control: Dynamic illusion effects in grasping?' Spatial Vision, 16: 3-4 pp.211-223
- Gallagher, S. (2012), 'The overextended mind', *Versus* 113 pp.57-68

- -   (2005), *How the Body Shapes the Mind,* Oxford: Oxford University Press

- -   (2000), 'Philosophical concepts of the self: Implications for cognitive sciences', *Trends in Cognitive Sciences* 4 pp.14–21
- -   (1986a), 'Body image and body schema: a conceptual clarification', *Journal of Mind and Behaviour* 7 pp.541–54
- -   (1986b), 'Hyletic experience and the lived body', *Husserl Studies* 3 pp.131-66
- Gallagher, S. and Chrisafi, A. (2009), 'Mental institutions', *Topoi* 28: 1 pp.:45-51
- Gallagher, S. and Cole, J. (1995) 'Body image and body schema in a deafferented subject', *Journal of Mind and Behaviour* 16, pp.369-390
- Gallagher, S. and Zahavi, D. (2007), *The Phenomenological Mind*, London: Routledge
- Galton, F. (1881), 'Visualised numerals', *Journal of the Anthropological Institute* 10 pp.85-102
- Gao, X., Xin, J., Sato, T., Hansuebsai, A., Scalzo. M., Kajiwara, K.,Guan, S., Valldeperas, J., Lis, M. and Billger, M. (2007), 'Analysis of cross-cultural colour emotion', *Colour Research and Application* 32:3 pp.223-9
- Gardner, M. (1970), 'Mathematical games: The fantastic combinations of John Conway's new solitaire game "life"', *Scientific American* 223 pp.120-3
- Gibson, J. (1986) [1979], *The Ecological Approach to Visual Perception,* NJ: Psychology Press
- Giummarra, M., Gibson, S., Georgiou-Karistianis, N. and Bradshaw, J. (2008), 'Mechanisms underlying embodiment, disembodiment and loss of embodiment', *Neuroscience and Biobehavioural Reviews* 32 pp.143-60
- -   (2007), 'Central mechanisms in phantom limb perception: The past, present and future, *Brain Research Reviews* 54 pp.219-32
- Glover, S. (2004), 'Separate visual representations in the planning and control of actions', *Behavioural and Brain Sciences,* 27 pp.3-78
- Goldin-Meadow, S. (2003), *Hearing Gesture: How our hands help us to Think,* MA: Harvard University Press
- Goldin-Meadow, S. and Singer, M. (2003), 'From children's hands to adults' ears: Gesture's role in the learning process', *Developmental Psychology* 39:3 pp.509-20
- Goller, A., Otten, L., and Ward, J. (2008) 'Seeing sound and hearing colours: An Event-related Potential Study of Auditory-Visual Synaesthesia', *Journal of Cognitive Neuroscience* 21:10 pp.189-81

- Gregory, R. (1980), 'Perceptions as hypotheses', *Philosophical Transactions of the Royal Society* B: 290 pp.181-97
- Grice, H. (1962), 'Some remarks about the senses' in R. Bulter ed. *Analytical Philosophy,* Oxford: Blackwell pp.133-53
- Grimes, J. (1996), 'On the failure to detect changes in scenes across saccades' in K. Atkins ed. *Vancouver Studies in Cognitive Science Vol. 5: Perception,* NY: Oxford University Press pp.89-110
- Gross, C. (1995), 'Aristotle on the brain', *The Neuroscientist* 1:4 pp.245-50
- Haines, R. (1991), 'A breakdown in simultaneous information processing' in G. Obrecht and L. Starks (eds.) *Presbyopia Research: From Molecular Biology to Visual Adaptation,* NY: Plenum Press pp.171-5
- Hatfield, G. (2002), 'Perception as unconscious inference' in D. Heyer and R. Mausfield eds. *Perception and the Physical World*, NY: John Wiley pp.115-42
- Head, H., and Holmes, G. (1911), 'Sensory disturbances from cerebral lesions', *Brain* 34 pp.102-254

- Heidegger, M. (2000) [1959-69], *Zollikon Seminars,* ed. M. Boss, trans. R. Askay and F. Mayr, IL: Northwestern University Press
- -   (1962) [1927], *Being and Time,* trans. J Macquarrie and E. Robinson, NY: Blackwell
- Hein, G., Lamm, C., Brodbeck, C., and Singer, T. (2011), ' Skin conductance response to the pain of others predicts later costly helping', *PLoS One* 6:8 e22759
- Held, R. and Hein, A. (1963), 'Movement-produced stimulation in the development of visually guided behaviour', *Journal of Comparative and Physiological Psychology* 56: 5 pp.872-6
- Helmholtz, H. (1962) [1860], *Handbook of Physiological Optics,* trans. J. Southall NY: Dover
- Hobson, A. (2009), 'The neurobiology of consciousness: Lucid dreaming wakes up', *International Journal of Dream Research* 2:2 pp.41-4
- Hoffding, S. (forthcoming), 'What is skilled coping? Experts on expertise', to appear in *Journal of Consciousness Studies*
- Hohwy, J. (2007), 'Functional integration and the mind', *Synthese* 159, pp.315-28
- Hollan, J., Hutchins, E. and Kirsh, D. (2000), 'Distributed cognition: toward a new foundation for human-computer interaction research', *ACM Transactions on Human-Computer Interaction* 7:2 pp.174-96
- Hume, D. (1977) [1777], *An Enquiry Concerning Human Understanding*, IN: Hackett
- Hurley, S. (2010), 'The varieties of externalism' in R. Menary ed. *The Extended Mind,* MA: MIT Press pp.101-54
- -   (1998), *Consciousness in Action,* MA: Harvard University Press
- Hurley, S. and Noë, A. (2003), 'Neural plasticity and consciousness', *Biology and Philosophy* 18 pp.131-68
- Husserl, E. (1999a), [1913], 'The structure of intentionality' (extract from *Ideas I*) in D. Welton ed. *The Essential Husserl,* IN: Indiana University Press pp.86-112
- -   (1999b), 'Horizons and the genesis of perception', trans. A. Steinbock in D. Welton ed. *The Essential Husserl*, IN: Indiana University Press pp.221-33
- -   (1991) [1893-1917], *On the Phenomenology of the Consciousness of Internal Time,* trans. J. Brough, Dordrecht: Springer

- - (1989), *Ideas Pertaining to A Pure Phenomenology and to a Phenomenological Philosophy: Second Book,* trans. R. Rojcewicz and A. Schuwer, Dordrecht: Kluwer
- - (1988) [1950], *Cartesian Meditations*, trans. D Cairns, Dordrecht: Kluwer
- - (1970) [1936], *The Crisis of the European Sciences and Transcendental Phenomenology,* trans. D. Carr, IL: Northwestern University Press
- Hutchins, E. (1996), *Cognition in the Wild,* MA: MIT Press
- Hutto, D. (2013), 'Exorcising action oriented representations: ridding cognitive science of its Nazgul', *Adaptive Behaviour* 21:3 pp.142-50
- Hutto, D. and Myin, E. (2013), *Radicalizing Enactivism*, MA: MIT Press
- Iriki, A., Tanaka, M. and Iwamura, Y. (1996), 'Coding of modified body schema during tool use by macaque postcentral neurones', *Neuroreport* 7:14 pp.2325-30
- Jackson, F. (2004), 'Representation and experience' in H. Clapin, P. Saines and P.Slezak, eds. *Representation in Mind,* London: Elsevier pp.107-24
- Jacob, P. (2006), 'Why visual experience is likely to resist being enacted', *Psyche* 12:1 online: http://www.theassc.org/files/assc/2630.pdf (last accessed 12/8/14)
- Jacob, P. and Jennerod, M. (2003), *Ways of Seeing: The Scope and Limits of Human Vision,* Oxford: Oxford University Press
- Jensen, R. (2009), 'Motor intentionality and the case of Schneider', *Phenomenology and the Cognitive Sciences* 8:3 pp.371-88
- Kaczmarek, K. Tyler, M. and Bach-y-Rita, P. (1997), 'Pattern identification on a fingertip-scanned electrotactile display', *Proceedings of the 19th IEEE/EMBS Conference* pp.1694-6
- Kant, I. (1977) [1783], *Prolegomena to any Future Metaphysics,* trans. J. Ellington, London: Hackett
- Keeley, B. (2002), 'Making senses of the senses: Individuating modalities in humans and other animals', *The Journal of Philosophy* 99:1 pp.5-28
- Kim J. (2006), *Philosophy of Mind,* CO: Westview Press
- Kirsh, D. (2009), 'Problem solving and situated cognition' in P. Robbins & M. Aydede eds. *The Cambridge Handbook of Situated Cognition,* Cambridge: Cambridge Press pp.264-30
- - (1995), 'The intelligent use of space', *Artificial Intelligence* 73: 1-2 pp.31-68
- Kirsh, D. and Maglio, P. (1994), 'On distinguishing epistemic from pragmatic action', *Cognitive Science* 18 pp.513-49
- Kiversten, J. and Farina, M. (2012), 'Do sensory substitution devices extend the conscious mind?' in F. Paglieri ed. *Consciousness in Interaction,* Amsterdam: John Benjamins pp.19-40
- Kiverstein, J., Farina, M. and Clark, A. (forthcoming), 'Substituting the senses' in M. Matthen ed. *The Oxford Handbook of the Philosophy of Perception,* Oxford: Oxford University Press (manuscript)
- Kiverstein, J. and Wheeler, M. eds. (2012), *Heidegger and Cognitive Science,* Basingstoke: Palgrave-Macmillan
- Kohler, I. (1962), 'Experiments with goggles', *Scientific American* 206 pp.299-309
- Krueger, J. (2014), 'Affordances and the musically extended mind', *Frontiers in Psychology*. 4: 1003:pp.1-12

- (2012), 'Seeing mind in action', *Phenomenology and the Cognitive Sciences*, 11.2:pp.149-173

- Krueger, J. and Overgaard, S. (2012), 'Seeing subjectivity: Defending a perceptual account of other mind',. *ProtoSociology* 47 pp.239-26

- Kyselo, M. and Walter, S. (2011), 'Belie integration in action: A defence of extended beliefs', *Philosophical Psychology* 24:2 pp.245-60

- LaBerge, S. (2004), *Lucid Dreaming*, CO: Sounds True Inc.

- (1990), 'Lucid dreaming: psychophysiological studies of consciousness during REM sleep' in R. Boozten , J. Kihlstrom and D. Schacter eds. *Sleep and Cognition,* Washington DC: American Psychological Association pp.109-26

- Langham, M., Hole, G., Edwards, J. and O'Neil, C. (2002), 'An analysis of 'looked but failed to see' accidents involving parked police vehicles', *Ergonomics* 45:3 pp.167-85

- Laureys, S., Pellas, F., Van Eeckhout, P., Ghorbel, S., Schnakers, C., Perrin, F., Berrê, J., Faymonville, M., Pantke, K., Damas, F., Lamy, M., Moonen, G. and Goldman, S. (2005), 'The locked-in syndrome: what is it like to be conscious but paralysed and voiceless?', *Progress in Brain Research* 150 pp. 495-511

- Lenay, C., Gappena, O., Henneton, S., Marque, C. and Genouelle, C. (2003), 'Sensory substitution: limits and perspectives', in Y. Hatwell, A.Steri and E. Gentaz eds. *Touching for Knowing, Cognitive Psychology of Haptic Manual Perception,* Amsterdam: John Benjamins pp.275-292

- Leon, M. (1988), 'Characterising the senses', *Mind & Language* 3:4 pp.243-70

- Levy-Tzedek, S., Hanassy, S., Abboud, S., Maidenbaum, S. and Amedi, A. (2012), 'Fast, accurate reaching movements with a visual-to-auditory sensory substitution device**,** *Restorative Neurology and Neuroscience* 30 pp.313-23

- Loughlin, V. (2013), 'Sketch this: Extended mind and consciousness extension', *Phenomenology and the Cognitive Sciences* 12:1 pp.41-50

- Loyd, D. (2007), 'Spatial limits on referred touch to an alien limb may reflect boundaries of visuo-tactile peripersonal space surrounding the hand', *Brain and Cognition* 64 pp.104-9

- Lurz, R. ed. (2009), *The Philosophy of Animal Minds*, Cambridge: Cambridge University Press

- Macpherson, F. (2010), 'Taxonimising the senses', *Philosophical Studies* 153 pp.123-42

- Maglio, P., Kandogan, E. and Haber, E. (2003), 'Distributed cognition and joint activity in collaborative problem solving', *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society,* online: http://www.almaden.ibm.com/coevolution/pdf/cogsci-cameraready-ver6.pdf (last accessed 22/8/14)

- Mainzer, K. (2007), *Thinking in Complexity,* NY: Springer

- Manzotti, R. (2011), 'The spread mind: Is consciousness situated?', *Teorema* 30:2 pp.55-78

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

- Maravita, A., Spence, C. and Driver, J. (2003), 'Multisensory Integration and the body schema: Close to hand and within reach', *Current Biology* 13 pp.531-9

- Marcel, A. (2003), 'The sense of agency: awareness of ownership of action' in J. Roessler and N. Eilan eds., *Agency and Self-Awareness,* Oxford: Oxford University Press pp.48-93
- Marr, D. (1982), *Vision*, SF: Freeman
- Martin, M. (2002), 'The transparency of experience', *Mind and Language* 4:4 pp.376-425
- \-   (1993), 'Sense modalities and spatial properties' in N. Eilan, R. McCarthy and B. Brewer (eds.) *Spatial Representation,* Oxford: Oxford University Press pp.206-17
- Maurer, D. and Maurer, C. (1988), *The World of the Newborn,* NY: Basic Books
- Maurer, D. and Mudlock, C. (2005), 'Neonatal synesthesia: A reevaluation' in L. Robertson and N. Sagiv eds. *Synaesthesia: Perspectives from Cognitive Neuroscience,* Oxford: Oxford University Press pp.193-213
- Maurer, D., Pathman, T., and Mondloch, C. (2006), 'The shape of boubas: sound-shape correspondences in toddler and adults', *Developmental Science* 9:3 pp. 316-22
- McCulloch, G. (2003), *The Life of the Mind,* Oxford: Oxford University Press
- McDonnell, P., Scott, R., Dickison, J., Theriault, R.and Wood, B. (1989), 'Do artificial limbs become part of the user? New evidence', *Journal of Rehabilitation Research and Development* 26:2 pp.17-24
- McGinn, C. (2004), *Mindsight,* MA: Harvard University Press
- McGurk H. and MacDonald, J. (1976) 'Hearing lips and seeing voices', *Nature* 264 pp.746-8
- McNeil, D. (2005), *Gesture and Thought*, Chicago: Chicago University Press
- Meijer, B. (1992), 'An experimental system for auditory image representations', *IEEE Transactions on Biomedical Engineering* 39:2 pp.112-21
- Menary, R. (2010a) ed. *The Extended Mind,* MA: MIT Press
- \-   (2010b) 'The holy grail of cognitivism: A Response to Adams and Aizawa', *Phenomenology and the Cognitive Sciences* 9 pp.605-18
- \-   (2010c), 'Cognitive integration and the extended mind' in R. Menary ed. *The Extended Mind,* MA:MIT Press pp.227-43
- \-   (2007a), *Cognitive Integration*, Basingstoke: Palgrave MacMillan
- \-   (2007b), 'Writing and thinking', *Language Sciences* 29 pp.621-32
- Merleau-Ponty, M. (2012) [1945], *Phenomenology of Perception,* trans. D. Landes, London: Routledge
- \-   (2004) [1945], 'Cezanne's doubt' in T. Baldwin ed. *Maurice Merleau-Ponty: Basic Writings,* London: Routledge pp.272-90
- \-   (1967) [1942] *The Structure of Behaviour,* trans. A. Fisher, Boston: Beacon Press
- \-   (1964a) [1961] 'Phenomenology and the sciences of man', trans. J. Wild in *The Primacy of Perception,* IL: Northwestern University Press pp.43-95
- \-   (1964b) [1947], 'The primacy of perception and its philosophical consequences', trans. J. Edie in *The Primacy of Perception,* IL: Northwestern University Press pp.12-42
- \-   (1964c) [1952], 'Indirect language and the voices of silence', in *Signs,* trans. R. McCleary, IL: Northwestern University Press pp.39-83
- \-   (1964d), 'The child's relations with others', trans. W. Cobb in *The Primacy of Perception,* IL: Northwestern University Press pp.96-155

- Metzinger, T. (2009), *The Ego Tunnel*, NY: Basic Books
- Miller, G. (1956), 'The magical number seven, plus or minus two: Some limits on our capacity for processing information', *Psychological Review* 63:2 pp.81-97
- Millikan, R. (1984). *Language, Thought, and other Biological Categories*, MA: MIT Press
- Milner, D. and Goodale, M. (2010), 'Two visual stream: Interconnections do not imply duplication of function', *Cognitive Neuroscience* 1:1 pp.65-8
- -   (2006), *The Visual Brain in Action,* Oxford: Oxford University Press
- -   (1992), 'Separate pathways for perception and action', *Trends in Neurosciences* 15:1 pp.20-5
- Mitroff, S., Simons, D., and Levin D., (2004), 'Nothing compares two views: Change blindness can occur despite preserved access to the changed information', *Perception & Psychophysics* 66 pp.1268-81
- Moguillansky, C., O'Regan, J., and Petitmengin, C., (2013), 'Exploring the subjective experience of the rubber hand illusion', *Frontiers in Human Neuroscience* 7:659 pp.1-18
- Mooney, T. (2011), 'Plasticity, motor intentionality and concrete movement in Merleau-Ponty', *Continental Philosophy Review* 44 pp.349-81
- Moore, G. E. (1953), *Some Main Problems of Philosophy,* London: George Allen and Unwin
- Moran, D. (1999), *Introduction to Phenomenology,* London: Routledge
- Morris, D. (2010), 'The enigma of reversibility and the genesis of sense in Merleau-Ponty', *Continental Philosophy Review* 43 pp.141-65
- Morris, J. and Twaddle, S. (2007), 'Anorexia Nervosa', *British Medical Journal* 334 pp.894-8
- Mosley, P. (2009), 'Bigorexia: bodybuilding and muscle dysmorphia', *European Eating Disorders Review* 17 pp.191-8
- Noë, A., (2012), *Varieties of Presence*, MA: Harvard University Press
- -   (2010), 'Vision without representation' in N. Gangopadhyay, M. Madary, and F. Spicer eds. *Perception, Action, and Consciousness*, Oxford: Oxford University Press pp.245-56
- -   (2008), 'Magic realism and the limits of intelligibility: What makes us conscious?' in J. Hawthorne ed. *The Philosophy of Mind*, Oxford: Blackwell pp.457-76
- -   (2004), *Action in Perception,* MA: MIT Press
- -   (2002) ed. *Is the Visual World a Grand Illusion?* Exeter: Imprint Academic
- O'Callaghan, C. (forthcoming), 'Not all perceptual experience is modality specific' in Stoke, Matthen and Biggs eds. (manuscript)
- Olsen, E. (2011), 'The extended self', *Minds and Machines* 21:4 pp.481-95
- O'Regan, J. (2014), 'The explanatory status of the sensorimotor approach to phenomenal consciousness, and its appeal to cognition' in Bishop and Martin (eds.), pp.23-35
- -   (2012), 'How to build a robot that is conscious and feels, *Minds and Machines* 22:2 pp.117-36
- -   (2011), *Why Red Doesn't Sound Like a Bell,* Oxford: Oxford University Press

- - (2010) 'Explaining what people say about qualia' in Gangopadhyay, Madary, and Spicer eds. eds. *Perception, Action and Consciousness,* Oxford: Oxford University Press pp.31-50

- O'Regan, J., Deubel, H., Clark, J., Rensinik, R. (2000), 'Picture changes during blinks: Looking without seeing and seeing without looking', *Visual Cognition* 7:1-3 pp.191-211

- O'Regan, J., and Myin, E. (2009), 'Situated perception and sensation in vision and other modalities' in. Robbins and Aydede eds. pp.185-200

- O'Regan, J., and Noë, A. (2001), 'A sensorimotor account of vision and visual consciousness', *Behavioural and Brain Sciences* 24 pp.939-1031

- Ortiz T, Poch J, Santos J., Requena C., Martı´nez A., Ortiz-Teran, L., Turrero, A., Barcia, J., Nogales, R., Calvo, A., Martinez, J., Cordoba, J., Pacual-Leone, A., (2011), 'Recruitment of occipital cortex during sensory substitution training linked to subjective experience of seeing in people with blindness', *PLoS One* 6:8 e23264

- O'Shaughnessy, B. (1995), 'Proprioception and the body image' in J. Bermudez, A. Marcel and N. Eilan eds. *The Body and the Self,* MA: MIT Press pp.175-203

- - (1980), *The Will: A Dual Aspect Theory Volume I,* Cambridge: Cambridge University Press

- Ou, L., Luo, M., Woodcock, A. and Wright, A. (2004) 'A study of colour emotion and colour preference part 1: Colour emotions for single colours', *Colour Research and Application* 29:3 pp.232-40

- Palmer, C., Paton, B., Hohwy, J and Enticott, P., (2013), 'Movement under uncertainty: the effects of the rubber hand illusion vary along the nonclinical autism spectrum', *Neuropsychologia* 51:10 pp.1942-51

- Patterson, J. and Grabois, M. (1986), 'Locked-in syndrome: a review of 139 cases', *Stroke* 17 pp.758-64

- Penfield, W. and Rassmusen, T. (1950), The Cerebral Cortex of Man: A Clinical Study of Localisation and Function, NY: MacMillan

- Pfeifer, R., Lungarella, M., Sporns, O., and Kuniyoshi, Y., (2007), 'On the information theoretic implications of embodiment –principles and methods', *Lecture Notes in Computer Science* 4850 pp.76-86

- Piaget, J. (1952), *The Origins of Intelligence in Children,* trans. M. Cook, NY: W.W. Norton & Co.

- Pisella, L., Sergio, L., Blangero, A., Torchin, H., Vighetto, A., Rossetti, Y. (2009), 'Optic ataxia and the function of the dorsal stream: Contributions to perception and action', *Neurophsychologia* 47 pp.3033-44

- Plato (2002), *Five Dialogues,* trans. G. Grube and J. Cooper, IN: Hackett

- Preston, B. (1998), 'Cognition and tool use', *Mind and Language* 13:4 pp.513-47

- Preston, B. and Bishop, M. (2002), *Views into the Chinese Room,* Oxford: Oxford University Press

- Preston, J.. (2010) 'The extended mind, the concept of belief, and epistemic credit' in Menary ed. *The Extended Mind,* MA: MIT Presspp.355-69

- Price, E. (2006), 'A critical review of congenital phantom limb cases and a developmental theory for the basis of body image', *Consciousness and Cognition* 15 pp.310-22

- Prinz., J. (2006), 'Putting the brakes on enactive perception', *Psyche* 12: 1, online: https://www.uoguelph.ca/~abailey/Resources/Prinz,%20Jesse%20--%20Putting%20the%20Brakes%20on%20Enactive%20Perception.pdf (last accessed 14/8/14)
- Proulx, M. and Stoerg, P. (2006), 'Seeing sounds and tingling tongues: Qualia in synaesthesia and sensory substitution', *Anthropology and Philosophy* 7:1-2 pp.135-50
- Putnam, H. (1975), 'The Meaning of Meaning', *Minnesota Studies in the Philosophy of Science* 7 pp.13193
- Ramachandran, V., and Blakeslee, S. (2005), *Phantoms in the Brain,* London: Harper Perennial
- Ramachandran, V. and Hubbard, E. (2003), 'Hearing colours, tasting shapes', *Scientific American* May 2003 pp.53-9
- - (2001a), 'Psychophysical investigations into the neural basis of synaesthesia, *Proceedings of the Royal Society of London B* 268 pp.979-83
- - (2001b), 'Synaesthesia: A window into perception, thought and language', *Journal of Consciousness Studies* 8:12 pp.3-34
- Ramachandran,V., Rogers-Ramachandran, D., and Cobb, S. (1995), 'Touching the phantom limb', *Nature* 337 pp.489-90
- Rao, R. and Ballard, D. (1999), 'Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects', *Nature Neuroscience Reviews* 2:1 pp.79-87
- Ratcliffe, M. (2007), *Rethinking Commonsense Psychology: A Critique of Folk Psychology, Theory of Mind and Simulation*, Basingstoke: Palgrave Macmillan
- Revonuso, R. (1999), 'Binding and the phenomenal unity of consciousness', *Consciousness and Cognition* 8 pp.173-85
- Roberts B., Harris M., Yates T., (2005), 'The roles of inducer size and distance in the Ebbinghaus illusion (Titchener circles)' *Perception* 34:7 pp.847 – 856
- Romdenh-Romluc, (2011a), *Routledge Philosophy Guidebook to Merleau-Ponty and Phenomenology of Perception,* London: Routledge
- - (2011b), 'Agency and Embodied Cognition', *Proceedings of the Aristotelian Society* 111 pp.79-95
- - (2007), 'Merleau-Ponty and the power to reckon with the possible' in T. Baldwin ed. *Reading Merleau-Ponty,* London: Routledge pp.44-58
- Rosenthal, D. (1997), 'A theory of consciousness' in N. Block, O. Flanagan and G. Guzeldere eds. *The Nature of Consciousness: Philosophical Debates,* MA: MIT Press pp.729-53
- - (1991) ed. *The Nature of Mind,* NY: Oxford University Press
- Ross, D. and Ladyman, J. (2010), 'The alleged coupling/constitution fallacy and the mature sciences' in Menary ed. pp.155-66
- Ross, D., Ladyman, J. and Kincaid, H. (2013) ed., *Scientific Metaphysics* Oxford: Oxford University Press
- Ross, P. (2001), 'Qualia and the senses', *The Philosophical Quarterly* 15:205 pp.495-511
- Rossetti, Y., Pisella, L. and Vighetto, L. (2003), 'Optic ataxia revisited: visually guided action versus immediate visuomotor control', *Experimental Brain Research* 153 pp.171-9
- Rowlands, M. (2010a), *The New Science of the Mind,* MA: MIT Press

- Rowlands, M. (2010b), 'Consciousness, broadly construed' in R. Menary ed. pp.271-94

-    (2003), *Externalism*, Buckinghamshire: Acumen

- Rudder-Baker, L. (2009), 'Persons and the extended mind thesis', *Zygon* 44: 3 pp.642-658
- Rupert, R. (2009), *Cognitive Systems and the Extended Mind,* Oxford: Oxford University Press

- Russell, B. (2002) [1912], *The Problems of Philosophy,* Oxford: Oxford University Press
- Ryle, G. (2009), 'Phenomenology versus 'The concept of mind'' in *Critical Essays: Collected Papers Volume I,* London: Routledge pp.186-204

-    (2000) [1949], *The Concept of Mind*, London: Penguin

- Sacks, O. (1995), *An Anthropologist on Mars: Seven Paradoxical Tales,* Basingstoke: Picador

-    (1985), *The Man Who Mistook His Wife for a Hat,* London: Picador
- Sartre, J-P. (2003) [1943], *Being and Nothingness,* trans. H. Barnes, London: Routledge
-    (2004) [1940], *The Imaginary*, trans. A. Eelkaim-Sartre and J. Webber, London: Routledge
- Saul, J.(forthcoming) 'Implicit bias, stereotype threat and women in philosophy' *to appear in F. Jenkins and K. Hutchison eds. Women in Philosophy: What Needs to Change?, Oxford: Oxford University Press*
- Schabowsky, C., Dromerick, A., Holley, S., Monroe, B., and Lum, P. (2008), ' Trans-radial upper extremity amputees are capable of adapting to a novel dynamic environment', *Experimental Brain Research* 188 pp.589-601

- Schellenberg, S. (2007), 'Action and self-location in perception', *Mind* 116:463 pp.603-32

- Searle, J. (1995), *The Construction of Social Reality,* London: Penguin
-    (1992), *The Rediscovery of the Mind,* MA: MIT Press
-    (1983), *Intentionality,* Cambridge: Cambridge University Press
-    (1980), 'Minds, brains and programs', *Behavioural and Brain Sciences* 3 pp.417-57
- Sellars, W. (1997) [1956], *Empiricism and the Philosophy of Mind,* MA: Harvard University Press
- Shenk, T. and McIntosh, R (2010), 'Do we have independent visual streams for perception and action?', *Cognitive Neuroscience* 1:1 pp.52-62

- Silverman, D. (2013), 'Sensorimotor enactivism and temporal experience', *Adaptive Behaviour*, forthcoming
- Simmer, J. and Ludwig, V. (2012), 'The colour of touch: A case of tactile-visual synaesthesia', *Neurocase: The Neural Basis of Cogniton* 18:2 pp.167-80
- Simmer, J., Mulvenna, C., Sagiv, N., Tsankanikos, E., Witherby, S., Fraser, C., Scott, K. and Ward, J. (2006), 'Synaesthesia: The prevalence of atypical cross-modal experiences', *Perception* 35 pp.1024-33

- Simons, D. and Chabris, C. (1999), 'Gorillas in our midst: Sustained inattentional blindness for dynamic events', *Perception* 28:9 pp.1059-74

- Simons, D. and Resnik, R. (2005), 'Change blindness: past, present, and future', *Trends in Cognitive Sciences* 9:1 pp.16-20

- Skinner, B. (1957), *Verbal Behaviour,* MA: Copley

- Smith, A. (2007), 'The flesh of perception: Merleau-Ponty and Husserl' in T.Baldwin ed. pp.1-22
- Smith, E. and Delargy, M. (2005), 'Locked-in syndrome', *British Medical Journal* 330 pp. 406-9

- Smith, Q. (1977), 'A phenomenological examination of Husserl's theory of hyletic data', *Philosophy Today* 2:4 pp.356-67

- Sorabj, R. (1971), 'Aristotle on demarcating the five senses', *Philosophical Review* 80:1 pp.55-79

- Spector, F. and Maurer, D. (2009), 'Synaesthesia: A new approach to understanding the development of perception', *Developmental Psychology* 45:1 pp.175-89

- Sprevak, M. (2009), 'Extended cognition and functionalism', *The Journal of Philosophy* 106 pp.503-21, manuscript from https://dl.dropboxusercontent.com/u/578710/homepage/Sprevak---Extended%20Cognition.pdf (last accessed 16/9/14)
- Spurrett, D. and Cowley, S., (2010), 'The extended infant: utterance activity and distributed cognition' in Menary ed. pp.295-324
- Stanley, J. and Williamson, T. (2001), 'Knowing how', *Journal of Philosophy* 98: 8 pp.411-44
- Steven, M. and Blakemore,C. (2004), 'Visual synaesthesia in the blind', *Perception* 37:7 pp.855-68
- Stiles, N., and Shimojo, S. (in press), 'Sensory substitution: a new perceptual experience' to appear in J. Wagemans ed. *The Oxford Handbook of Perceptual Organisation,* Oxford: Oxford University Press (manuscript)
- Stokes, D. and Biggs, S. (forthcoming), 'The dominance of the visual', to appear in D. Stokes, M. Matthen and S. Biggs eds. *Perception and its Modalities,* Oxford: Oxford University Press
- Stokes, D., Matthen, M. and Biggs, S. eds. (forthcoming) *Perception and its Modalities*, Oxford: Oxford University Press
- Stoneham, T. (2008), 'A neglected account of perception', *Dialectica* 62:3 pp.307-2
- Stratton, G.(1897), 'Vision without inversion of the retinal image', *The Psychological Review* 4:4 pp.341-60
- Strawson, G. (2006), 'Realistic monism: Why physicalism entails panpsychism', *Journal of Consciousness Studies* 13:10-11 pp.184-280
- Sutton, J., Harris, C., Keil, P., and Barnier, A. (2010), 'The psychology of memory, extended cognition, and socially distributed remembering', *Phenomenology and the Cognitive Sciences*, 9:4 pp.521-560
- Takahashi, C., Diedrichsen, J. and S. Watt (2009), 'Integration of vision and haptics during tool use', *Journal of Vision* 9:6 pp.1-13

- Thompson, E. (2007), *Mind in Life,* MA: MIT Press
- Tong, F. (2003), 'Primary visual cortex and visual awareness', *Nature Reviews: Neuroscience* 4 pp.219-29
- Tsakiris, M., and Haggard, P. (2005), 'Experimenting with the acting self', *Cognitive Neuropsychology* 22:3-4 pp.387-407
- Tsakiris, M., Carpenter, L., James, D., and Fotopoulou, A., (2010), 'Hand only illusion: multisensory integration elicits sense of ownership for body parts but not for non-corporeal objects', *Experimental Brain Research* 204 pp.343-52
- Tsakiris, M., Prabhu, G., and Haggard, P. (2006), 'Having a body versus moving your body: How agency structures body-ownership', *Consciousness and Cognition* 15:2 pp.423–32
- Tskiris, M., Schutz-Bosbach, S. and Gallagher, S. (2007), 'On agency and body ownership: Phenomenological and neurocognitive reflections', *Consciousness and Cognition* 16 pp.645-60
- Turner, J. (forthcoming), 'Super organisms and superindividuality: the emergence of individuality in a social insect assemblage' in F. Bouchard & P. Hunemann eds. *From Groups to Individuals*, MA: MIT Press
- Tye, M. (2002), 'Qualia and visual content revisited' in D. Chalmers ed. *Philosophy of Mind: Classical and Contemporary Readings*, Oxford: Oxford University Press
- Van Gelder, T. (1998), 'The dynamical hypothesis in cognitive science', *Behavioural and Brain Sciences* 21:5 pp.615-28
- Varela, F., Thompson, E. and Rosch, E., (1991) *The Embodied Mind: Cognitive Science and Human Experience,* MA: MIT Press
- Villegas, J., Castro, E., and Gutierrez, J. (2009), 'Representations in problem solving: a case study with optimization problems', *Electronic Journal of Research in Educational Psychology* 17:7 pp.279-308
- Voss, U., Frenzel, C., Koppehele-Gossel, J. and Hobson, A. (2012), 'Lucid dreaming: an age-dependent brain disassociation', *Journal of Sleep Research,* forthcoming
- Ward, D. (2012), 'Enjoying the spread: Conscious externalism reconsidered', *Mind* 121:483 pp.731-51
- Ward, J. and Wright, T. (2012) 'Sensory substitution as an artificially acquired synaesthesia', *Neuroscience and Bio Behavioural Reviews* (in press) (manuscript)
- Ward, J.and Simmer J. (2005), 'Is synaesthesia an x-linked dominant trait with lethality in males?', *Perception* 34 pp.611-23
- Watson, J. (1913), 'Psychology as the behaviourist views it', *Psychological Review* 20 pp.158-77
- Weiskopf, D. (2008), 'Patrolling the mind's boundaries' *Erkenntnis* 68 pp.265–276.
- Westbury, H. and Neumann, D. (2008), 'Empathy-related responses to moving film stimuli depicting human and non-human animal targets in negative circumstances', *Biological Psychology* 78:1 pp.66-74
- Wheeler, M. (2013) 'Is cognition embedded or extended? The case of gesture' in Z. Radman ed. *The Hand, an Organ of the Mind* MA: MIT Press pp.269–301
- -   (2010), 'In defence of extended functionalism', in Menary ed. *The Extended Mind* MA: MIT Press pp.245-70
- -   (2007) *Reconstructing the Cognitive World,* MA: MIT Press

- Wilson, R. (2004), *Boundaries of the Mind: The Individual in the Fragile Sciences,* Cambridge: Cambridge University Press
- Wilson, R. (2010), 'Extended vision' in Gangopadhyay, Madary and Spicer eds. pp. 271-90
- - (2004), *Boundaries of the Mind: The Individual in the Fragile Sciences*, NY: Cambridge University Press
- (1994), 'Wide computationalism", Mind 103 pp.351-72
- Wittgenstein, L. (1958), *Philosophical Investigations,* trans. G. Anscombe, NJ: Prentice Hall
- Yarbus, A. (1967), *Eye Movements and Vision,* trans. B. Haigh, NY: Plenum Press
- Zahavi, D. (2003), *Husserl's Phenomenology*, Stanford: Stanford University Press

Web Resources

- 'The man who lost his body', first broadcast October 16th 1997, BBC2 (UK), https://www.youtube.com/watch?v=AvLRwRAAoww (last accessed 31/8/14)
- 'Authors@google: Alva Noë', https://www.youtube.com/watch?v=af3Vq-C1ck8
- 'Selective attention test', https://www.youtube.com/watch?v=vJG698U2Mvo (last accessed 8/8/14)
- 'The vOICe', www.seeingwithsound.com (last accessed 14/7/14).
- 'Soda can collecting robot' [Herbert demo], https://www.youtube.com/watch?v=YtNKuwiVYm0 (last accessed 22/8/14)
- 'Shakey: Experimentation in robot learning and planning (1969)', https://www.youtube.com/watch?v=RhrLHkVuerc (last accessed 22/8/14)