Studies in Applied Philosophy, Epistemology and Rational Ethics

SAPERE

John Mark Bishop
Andrew Owen Martin *Editors*

Contemporary Sensorimotor Theory

Springer

Studies in Applied Philosophy, Epistemology and Rational Ethics

Volume 15

For further volumes:
http://www.springer.com/series/10087

About this Series

Studies in Applied Philosophy, Epistemology and Rational Ethics (SAPERE) publishes new developments and advances in all the fields of philosophy, epistemology, and ethics, bringing them together with a cluster of scientific disciplines and technological outcomes: from computer science to life sciences, from economics, law, and education to engineering, logic, and mathematics, from medicine to physics, human sciences, and politics. It aims at covering all the challenging philosophical and ethical themes of contemporary society, making them appropriately applicable to contemporary theoretical, methodological, and practical problems, impasses, controversies, and conflicts. The series includes monographs, lecture notes, selected contributions from specialized conferences and workshops as well as selected PhD theses.

Advisory Board

John Mark Bishop · Andrew Owen Martin
Editors

Contemporary Sensorimotor Theory

Springer

Editors
John Mark Bishop
Department of Computing
Goldsmiths College
University of London
London
United Kingdom

Andrew Owen Martin
Department of Computing
Goldsmiths College
University of London
London
United Kingdom

ISSN 2192-6255 ISSN 2192-6263 (electronic)
ISBN 978-3-319-05106-2 ISBN 978-3-319-05107-9 (eBook)
DOI 10.1007/978-3-319-05107-9
Springer Cham Heidelberg New York Dordrecht London

Library of Congress Control Number: 2014931506

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

For the enchantress of numbers

Preface

This book emerged following the first AISB member's workshop on 'Sensorimotor Theories of Perception', held at Goldsmiths, University of London on 26th Sep 2012. At the workshop contributors were asked to comment on the contemporary state of sensorimotor theories of perception and reflect on their place in modern Cognitive Science. As editors and workshop convenors, we would like to thank all the workshop participants for making the event so thought provoking and enjoyable and particularly those who continued engagement with the project by contributing to this volume.

Contents

Contemporary Sensorimotor Theory: A Brief Introduction

J. Mark Bishop and Andrew O. Martin

Dept. Computing, Goldsmiths, University of London, UK
m.bishop@gold.ac.uk

1 Background

'Sensorimotor Theory' offers a new *enactive approach*[1] to perception that emphasises the role of motor actions and their effect on sensory stimuli. The seminal publication that launched the field is the target paper co-authored by J. Kevin O'Regan and Alva Noë and published in Behavioral and Brain Sciences (BBS) for open peer commentary in 2001 [27].

In the central argument of their paper, O'Regan and Noë suggest radically shifting the nexus of research in visual perception away from analysis of the raw visual patterns of stimulation, to refocus on the law-like changes in visual stimulation brought about as a result of an agent's actions in the [light-filled] world.

A key consequence of this change is a new way of characterising objects by the unique set of 'sensorimotor correspondences' that define the characteristic changes in objective appearance brought about by the agent-object interactions [in the world]. These characteristic correspondences - relating the movement of any object relative to the agent - define its *sensorimotor dependencies* [qua world]; an agents practical knowledge of these sensorimotor dependencies constitutes its visual experience.

Thus in O'Regan and Noë's sensorimotor theory, perhaps for the first time, we have a rich, testable, psychological (and philosophically grounded) theory that accounts for *why* our conscious experience of the world appears as it does. This is a significant achievement and one that, in our opinion, goes a long way to answering at least some of the *hard problems of consciousness*[2].

This shift of nexus was, in part, informed by experiments in visual perception performed by O'Regan; for example, O'Regan, Rensink and Clark's discovery

[1] The term 'enactive approach' is taken from Noë [25] where he states, "*What I call here the enactive approach was first presented in [27]. We refers to the view as the sensorimotor contingency theory. Hurley and I, in joint work, deploy another term: the dynamic sensorimotor account. I borrow the term enactive from Francisco Varela and Evan Thompson (Varela, Thompson and and Roesch 1991 [49]), although I may not use it in exactly their sense. I use the term because it is apt, and to draw attention to the kinship of our view and theirs.*"

[2] David Chalmers introduced the term 'hard problem' to investigate, "Why is all this [neural] processing accompanied by an inner life?" [12]; we deploy the phrase 'hard problems of consciousness' to additionally encompass related problems pertaining to Levine's 'explanatory gap' [21].

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_1,

of *change blindness*[3] presents a challenge to the classical characterisation of perception as resulting in the construction of a rich, hi-fidelity mental representation[4]. To account for such experimental data (and many other peculiarities of [visual] consciousness[5]) a new conception of visual perception was outlined by O'Regan and Noë in their BBS article *A sensorimotor account of vision and visual consciousness*:

> .. the central idea of our new approach is that *vision is a mode of exploration of the world that is mediated by knowledge of what we call* ***sensorimotor contingencies***[6].

Thus O'Regan and Noë's sensorimotor framework for perception shifted the problem of vision away from that of construction of rich internal representations of an 'out there' world to that of active exploration of the environment 'on demand'; conscious experience being *brought forth* via a series of [saccadic] movements that either confirm (or disabuse) the notion that the world *actually is* of the form currently anticipated[7]. In this view the world is effectively used as its own external memory and the 'objects of the world' as their own 'representation'; hence removing at a stroke the absolute requirement to store and update a rich, hi-fidelity internal representation of the rich, out-there world.

This concept of perception as an 'active interrogation' (i.e. the exploration of external visual features as and when they are required) contra 'passive snapshot' (and subsequent rumination over rich internal mental representations) naturally accounts for why some perceptual changes might go unnoticed - *even if they are specifically attended to*[8] - if they are not in conflict with the agent's [current] anticipation(s).

---

[3] .. where significant portions of a subject's visual field can change and yet not be consciously experienced due either to the transition being gradual, or the subject's attention being distracted by other simultaneous changes; e.g. by the appearance so called 'mud-splat' distractors across the visual field.

[4] Problems arise in the classical view because, for example, large visual changes would naturally be anticipated to be perceived if there were large-enough differences between the state of the environment and the state of the internal representation.

[5] For example, the extra-retinal signal; trans-saccadic fusion; saccadic suppression; 'filling in' of the blind spot and other retinal non-homogeneities.

[6] In this introduction, after Broackes [27], we prefer the term 'sensorimotor dependencies' rather than 'sensorimotor contingencies', as the latter "has perhaps unfortunate connotations of non-necessity".

[7] O'Regan [29] reports the neuroscientist Donald MacKay as asserting, "*the eye was like a giant hand that samples the outside world*".

[8] For example, by slowly changing the colour of a merry-go-round, O'Regan and team demonstrate how large-scale changes to a visual scene can go unnoticed by even the most attentive of perceivers. NB. At first sight the phenomena of change blindness seems similar to the phenomenon of 'inattentional blindness', where something that is fully in view is not noticed because attention is elsewhere (cf. Daniel Simons' now ubiquitous invisible 'Gorilla in our midst' experiment); but change blindness is conceptually a very different effect since it depends crucially on the occurrence of a brief transitory event in the visual field that *distracts* attention, (instead of simply depending on the fact that attention is elsewhere).

In their response to the BBS peer commentaries [27] O'Regan and Noë particularly emphasize the twin notions of 'bodiliness' and 'grabbiness'. These are, respectively, features of the worldly stimulation that refer to the way in which sensory stimulation changes as bodily actions are performed, and the [involuntary] power of some sensory stimulation(s) to grab the entity's 'attention'.

O'Regan and Noë develop these notions to (a) help distinguish between perceived [real] and conceived [imagined] experience and (b) help illustrate how the sensorimotor conception of visual perception dissolves various degrees of what in 1983 Levine [21] termed 'the explanatory gap' - *the gulf that separates physical processes in the brain from the experienced quality of sensations.*

In their subsequent BBS discussion O'Regan and Noë focus on three different aspects to this 'gap':

Intra-modal: why does an object look *like this* (e.g. spherical), rather than *like that* (e.g. cubical)?
Inter-modal: what makes one experience *visual* whereas another is *tactile*?
Absolute: why is there *any* subjective experience at all?

.. and claim that sensorimotor theory makes a contribution to understanding (and potentially closing) all three of these aspects of the explanatory gap; thus:

- sensorimotor theory helps close the intra-modal gap because the differences between, say, a spherical and a cubical object of sight correspond in part to differences in degrees in which the specific laws of [visual] sensorimotor contingency are exhibited;
- sensorimotor theory helps close the inter-modal gap because of the different types of sensorimotor dependencies at play across the different sensory modalities. For example, 'moving your eyes to the left or right will produce a change in sensory stimulation related to an object if that object is being visually perceived; but not if it is being tactually perceived, or if it is being listened to'.
- however consideration of the third, most fundamental of the explanatory gaps, reveals subtle differences in the way in which Kevin and Alva conceive the extent of their theory ..

To illustrate the latter point of departure between O'Regan and Noë we must first consider an additional 'explanatory gap' - one that sits between the inter-modal and the absolute - and which emphasises the 'perceptual aspect': what is the basis of the difference between perceptual and non-perceptual awareness of a thing? For example O'Regan and Noë suggest to compare the experience of seeing a book on a table in front of you with that of your [non-perceptual] awareness of a book on a bookshelf in the room next door; and describe how two features of the experience serve to disambiguate the two situations:

Bodiliness: movement(s) of your body affect your perception of the book in front of you *but not* your awareness of the book in the room next door.

Grabbiness: sudden changes in the visual field (say a flashing light) when looking at the book in front of you provoke an ‘automatic orientating reflex’ that cause eyes to saccade towards the light; such visual *grabbiness* is entirely absent from your experience of the book on the bookshelf.

As O’Regan and Noë summarise:

> Bodiliness and grabbiness explain not only the difference between seeing something and merely thinking about it [..] but they also explain why seeing something has the sort of qualitative characteristics that it does.

In later writings [for example herein] O’Regan also considers a third feature of the real world, its ‘insubordinateness’ - how aspects of the world can change outside of the agent’s command - and contrasts the ‘insubordinateness’ of the real world view of the book in front of the agent, with the [virtual] perception of the book on the shelf in the adjacent room (where no such effects pertain)[9].

In *Why red doesn’t sound like a bell* [29] Kevin further extends the scope of bodiliness and grabbiness by conceiving them as two dimensions of a ‘phenomenality plot’; a graph which offers insight into the age-old philosophical question of why objects are phenomenally perceived in the way they are[10].

The phenomenality plot positions various experiences (e.g. hunger; an itch; driving etc.) on a two-dimensional graph, with the x-axis defining increasing bodiliness and the y-axis increasing grabbiness, such that *mental states that are located high up on the main diagonal, that is, states that objectively possess high bodiliness and high grabbiness, are precisely those that people will tend to say have a real sensory feel; that is, they have high “phenomenality”, or “something it’s like”*, (ibid).

Thus, where O’Regan [28] and Noë [26] agree [vis a vis sensorimotor theory closing the absolute explanatory gap] is that, after Nagel [24], there is “something it is like” to have a sensory experience and that sensorimotor theory can go a long way in (a) characterising what this likeness is, and (b) shedding insight on the nature of phenomenal perception itself.

---

[9] In [29] Kevin also discusses a fourth aspect of sensory interactions in the real world that could explain the presence of raw feel and which helps to distinguish real experience from imagined; its ‘richness’, however this fourth aspect is absent from Kevin’s most recent accounts of sensorimotor theory.

[10] “*The sensorimotor approach to raw feel explains four mysteries: why there’s something it’s like to have a raw sensory feel (presence); why different raw feels feel different; why there is structure in the differences; and why raw feel is ineffable. Of these four mysteries, the one philosophers consider to be the most mysterious is the first one, namely the question of why there’s ‘something it’s like’ to have an experience. If richness, bodiliness, insubordinateness, and grabbiness are the basis for the what it’s like of sensory feels, then we can naturally ask whether these concepts can do more work for us. In particular we would like them to explain, in the case of other, nonsensory types of experiences, the extent to which people will claim these have a ‘something it’s like’* ”, O’Regan [29].

> First, visual experience is visual, rather than, say, tactual or olfactory. Second, it is forcibly present to your consciousness. Third, it is ongoing, that is, the experience seems to be happening to you in a continuous way: its subjective character lasts while the experience continues. Fourth, the experience strikes us as ineffable, that is, though you experience it as possessing various qualities, the exact qualitative character escapes description in words.
>
> We believe the sensorimotor approach allows us to explain each of these aspects of the quality of the experience. To the extent, then, that the experience itself is constituted by the presence of just these qualities, then the sensorimotor account can explain why the experience occurs at all.

Where O'Regan and Noë begin to part company is in dealing with the apparent promiscuity of the sensorimotor account; a tacit implication of which being that *any medium of interaction with the environment will produce a conscious experience in the agent* with the exact profile of the conscious experience merely contingent upon the profile of sensorimotor dependencies. As Clark and Toribio wryly observed in their response to O'Regan and Noë's magnum opus:

> A good ping-pong playing robot, which uses visual input, learns about its own sensorimotor contingencies, and puts this knowledge to use in the service of simple goals (e.g., to win, but not by too many points) would meet all the constraints laid out. Yet it seems implausible to depict such a robot (and they do exist see, e.g., Andersson 1988 [2]) as enjoying even some kind of modest visual experience. Surely someone could accept all that O&N offer, but treat it simply as an account of how certain visual experiences get their contents, rather than as a dissolution of the so-called hard problem of visual qualia.

In their joint response to Clark and Toribo's examination of the internal life of a ping-pong playing robot O'Regan and Noë simply retort that:

> .. as it is described, it is simply far too simple to be a plausible candidate for perceptual consciousness of the kind usually attributed to animals or humans [...] once we imagine a robot that not only masters sensorimotor contingencies, but makes use of that mastery to engage with the world in a thoughtful and adaptable way, it becomes necessary to say that it has (at least primitive) visual experiences.

However, in later writings Noë appears to retreat from this position a little; for example in *Action and Perception* [25] he highlights both that:

> Nothing in our view committed us to saying that the robot would be perceptually conscious. All we committed ourselves to is the possibility that the robot could be perceptually conscious if it acquired the relevant practical skills, (ibid., chapter 7, footnote 12).

and emphasises, after Thompson [46], a fundamental link between 'mind and life'..

> .. with increasing sensorimotor complexity you get the appearance of a life-form that embodies a measure of sensitivity to the way its own movements change the way the environment stimulates it. In this way, and with a healthy dose of handwaving, we make plausible the idea that the emergence of perceptual consciousness in the biological world is, in effect, a matter of the emergence of cognitive agents with sensorimotor capabilities. There aren't two stories: the saga of the emergence of cognitive agents, and then that of the appearance of consciousness. Consciousness and cognition are themselves aspects of the development of life", (ibid., pp. 230)[11].

In contrast, in O'Regan's later writings (and specifically in this volume), to distinguish conscious agents from purely reactive systems (e.g. ping-pong playing robots; missile guidance systems etc) Kevin simply further refines precisely what is required of a [more sophisticated] sensorimotor agent in order for it to genuinely instantiate phenomenal consciousness, by appealing to ..

> .. a particular, hierarchical, form of *cognitive access* similar to that used in the higher-order theories (HOT) of consciousness [38] [10]".

As O'Regan states (ibid):

> The trick used in the sensorimotor approach is to try to 'dissolve' the hard problem of qualia rather than 'solving' it .. [that is, we need to] .. decompose consciousness into two parts: the sensorimotor interactions whose laws constitute experienced quality; and a form of cognitive access which makes that quality conscious.

Thus, for an agent to have the conscious experience of a quality, *say redness*, the agent must - via its interaction with the environment - both instantiate appropriate:

sensorimotor dependencies appropriate to pertaining to *redness*;

cognitive access to the actions being performed, such that the agent may claim, 'I am doing this'.
Such 'cognitive access' introduces additional requirements for the agent to have:
1. *a notion of a self* and thus have knowledge about its own body, mind and social context;
2. *access* to the 'experienced' quality.

[11] This point is further finessed in private communication (email: 08/10/13) in which Alva further clarified, "I am inclined to think that if you start out with a spark of consciousness, the sensorimotor enactive account can explain how you end up with all the varieties of consciousness. But sms alone won't tell you how we get started. In my view - and here I do agree with Thompson - you need life to get the ball rolling."

That is, the agent must exhibit the right sensorimotor dependencies and also (i) be poised to make use of its phenomenal knowledge [of the experiential quality] and (ii) be 'aware' that it is poised to make use of this knowledge[12].

As O'Regan summarises [in this volume]:

> When I say it feels like something rather than nothing to have a sensory experience, this statement can be decomposed into being conscious of facts like: it's modal (because it has the properties of touch, not of other senses), it has a particular quality (determined by the particular sensorimotor law), it's real (because it has the bodiliness, insubordinateness and grabbiness of real interactions, and furthermore controls my thoughts the way real things do), etc. The claim is that there is nothing more to phenomenal consciousness than potential conscious access to all these facts.

Thus, in an attempt to close the 'absolute gap' via sensorimotor theory it seems one is obliged to follow either Noë and reify a link between mind and life, or tread in O'Regan's footsteps and insist on the need for additional explicit cognitive access to phenomenal consciousness (and in the process potentially over-complicating, and hence losing some of the elegance of, foundational sensorimotor theory).

However, even aside from its extra complexity, a key implication of O'Regan's sensorimotor theory (but not Noë's) is to abstract 'consciousness' from any causal connection with its material substrate; as long as the putative conscious agent exercises the appropriate sensorimotor contingencies and concomitant 'cognitive access', phenomenal consciousness 'drops out' for free. If correct, this conceptual finding is most helpful to those roboteers who dream of building a conscious robot; and such pioneers are no straw-men..

For example, as early as 2002 Kevin Warwick (from the University of Reading, UK) stated that he believed machine consciousness was old news; his so called 'seven dwarf robots' already "instantiate a machine consciousness, albeit in a very weak form. They are, perhaps, if a comparison is to be drawn in terms of complexity, as conscious as a slug", [50].

This optimistic view of machine consciousness was subsequently echoed at the very highest stratum of UK research when, in 2004, Owen Holland (from the University of Essex) led a team which won nearly half a million pounds of UK research council funding to realise the goal of instantiating machine consciousness in a humanoid robot called Cronos (through appropriate internal computational modelling of robot-self and robot-external-world)[13].

---

12 As an example consider an agent being poised to press a car brake as it approaches a red traffic light and being aware that it is poised to press the break pedal.

13 In their EPSRC funding application Owen's team wrote that they expected "to enable some of the robots to be regarded as possessing a form of machine consciousness".

And these dreams of machine consciousness are closely aligned to Kevin's own vision of future robotics[14]; for example in chapter 6 of 'Why red doesn't sound like a bell' he remarks:

> When I was a child, my dream was to build a conscious robot. At that time in the 1960s, computers were just coming into existence, and they were where I hoped to find inspiration";

Furthermore, in chapter 7 (ibid) Kevin concludes:

> Where are we now with respect of my dream of building a conscious robot? So far I have been trying, one by one, to lay aside aspects of consciousness that pose no logical problem for science and that could be - and to a degree, in a few instances, are now - attainable in robots. I started with thought, perception, and communication, and went on to the self. I hope to have shown that even if we are still a long way from attaining human levels of these capacities in a robot, there is no logical problem involved in doing so.

However the formalisation and abstraction of consciousness required to build any such 'tin-can robot'[15] also implies that the instantiation of consciousness is not strongly contingent on the precise physical instantiation of theory; for example, the same conscious experience must arise if systems were instantiated by a [suitably rich] virtual reality (VR) computational simulation of a particular mode of [virtually] 'embodied interactions'.

Thus if one of Kevin Warwick's real 'seven dwarf' robots is as conscious as a slug as it flits about its coral, then a VR simulation of that robot would be as conscious as a slug as it flits about its virtual coral; if Owen Holland's real-world Cronos instantiates a form of machine consciousness as it moves a pint-glass on a table, then a VR simulation of Cronos must instantiate a similar form of machine consciousness as it moves a virtual pint-glass about a virtual table; and if a real robot constructed along O'Regan's sensorimotor principles experienced the particular phenomenal sensation of the 'squidginess of a sponge' [as it squeezed a sponge] then a VR simulation of that robot enacting the same

---

[14] A summary of Kevin's positive position vis a vis machine consciousness as espoused herein is also presented in: *Why red doesn't sound like a bell* [29]; *How to make a robot that feels*, the talk he gave at the PT-AI Conference in Thessaloniki Greece in 2011, and most recently in the paper *How to Build a Robot that is Conscious and Feels* [30].

[15] A 'tin-can robot' is any robotic device where the physical material of its embodiment is relatively unimportant; cf. Tom Ziemke's demarcation of embodiment into strong and weak forms. In a weakly embodied cognitive system the actual material of embodiment doesn't affect the cognitive states of the system and is important only in so far as it enables the system to appropriately support various high level functions (e.g. to be functional a robot arm must at least be strong enough to move itself); in a strongly embodied system, the actual material of embodiment does matter in its giving rise to genuine cognitive states.

sensorimotor dependencies [and cognitive access] of squidging a virtual sponge must experience the same phenomenal quality of 'squidginess'.

Thus a corollary of O'Regan's move to close the absolute gap [by a refinement of merely abstract formal processes] implies that it - in common with the other two "VR robot systems" highlighted above - *would be phenomenally conscious purely in virtue of its execution of an appropriate computer program*; and hence that it would be vulnerable to the various critiques of computationalism [7], in so far as these critiques hold at all. E.g.

1. That computation cannot give rise to semantics (cf. Searle's Chinese room argument [39]).
2. That computation cannot give rise to mathematical insight (cf. Penrose's Gödelian argument [32]).
3. That computation cannot give rise to consciousness (cf. Bishop's 'dancing pixies' argument [5]).

1.1 The Chinese Room Argument

Perhaps the most well known critic of computational theories of mind is John Searle. His best-known work on machine understanding, first presented in the 1980 paper 'Minds, Brains & Programs' [39], has become known as the Chinese room argument (CRA). The central claim of the CRA is that computations alone are not sufficient to give rise to cognitive states, and hence that computational theories of mind cannot fully explain human cognition. More formally Searle stated that the CRA was an attempt to prove the truth of the premise:

- Syntax is not sufficient for semantics;

which, together with the following two axioms ..

- Programs are formal (syntactical);
- Minds have semantics (mental content);

... led him to conclude that 'programs are not minds' and hence that computationalism - the idea that the essence of thinking lies in computational processes and that such processes thereby underlie and explain conscious thinking - is false [42].

It is beyond the scope of this introduction to summarise the extensive literature on the CRA[16]; however in *A view into the Chinese room* Bishop [6] summarised Searle's core argument as follows:

> .. Searle describes a situation whereby he is locked in a room and presented with a large batch of papers covered with Chinese writing that he does not understand. Indeed, Searle does not even recognise the symbols

[16] For a broad selection of essays detailing these and other critical arguments see Preston and Bishop's edited collection '*Views into the Chinese room*' [34].

> as being Chinese, as distinct from say Japanese or simply meaningless patterns. Later Searle is given a second batch of Chinese symbols, together with a set of rules (in English) that describe an effective method (algorithm) for correlating the second batch with the first, purely by their form or shape. Finally Searle is given a third batch of Chinese symbols together with another set of rules (in English) to enable him to correlate the third batch with the first two, and these rules instruct him how to return certain sets of shapes (Chinese symbols) in response to certain symbols given in the third batch.
>
> Unknown to Searle, the people outside the room call the first batch of Chinese symbols 'the script', the second set 'the story', the third 'questions about the story' and the symbols he returns they call 'answers to the questions about the story'. The set of rules he is obeying they call 'the program'.
>
> To complicate matters further, the people outside the room also give Searle stories in English and ask him questions about these stories in English, to which he can reply in English.
>
> After a while Searle gets so good at following the instructions and "outsiders" get so good at supplying the rules he has to follow, that the answers he gives to the questions in Chinese symbols become indistinguishable from those a real Chinese person might give.
>
> From an external point of view, the answers to the two sets of questions, one in English the other in Chinese, are equally good; Searle, in the Chinese room, has passed the Turing test. Yet in the Chinese language case, Searle behaves 'like a computer' and does not understand either the questions he is given or the answers he returns, whereas in the English case, ex hypothesi, he does. Searle contrasts the claim posed by members of the AI community - that any machine capable of following such instructions can genuinely understand the story, the questions and answers - with his own continuing inability to understand a word of Chinese..

However, the thirty plus years since its inception have witnessed many reactions to the Chinese room argument - ranging across communities in cognitive science, artificial intelligence, linguistics, anthropology, philosophy and psychology - with perhaps the most widely held criticism of Searle's position being based on what has become known as the 'Systems Reply'. This concedes that, although the person in the room doesn't understand Chinese, the entire system (of the room, the person and its contents) does.

Not surprisingly Searle finds this response entirely unsatisfactory and responds by allowing the person in the room to internalise everything (the rules, the batches of paper etc) so that there is nothing in the system not internalised within Searle. Now in response to the questions in Chinese and English there are two subsystems, a native monoglot English speaking Searle and an apparently Chinese fluent Searle, busy internalising the Chinese room. All the same he [John Searle] trenchantly continues to insist that he understands nothing of

Chinese, and a fortiori, neither does the internalised system because there isn't anything in the internalised system that is not just a part of him.

Thus, if Searle is told a joke in English, he will laugh and enjoy the experience of finding the joke funny; but if he is told a joke in Chinese, even if his 'internalised Chinese room' dictates he outputs the appropriate 'Chinese ideograph(s) for amusement', he will never experience concomitant laughter as, even with the Chinese room internalised, he cannot ever *get* the joke ..

The fiercesome power of Searle's Chinese room argument is such that, if correct, it demonstrates that no computational systems will 'ever genuinely understand' [Chinese] and hence that *all computational explanations must ultimately fail to provide an adequate model for cognition*; as Michael Wheeler presciently observes [51]:

> .. the extent that the Chinese Room argument succeeds in burying the idea that computation is sufficient for mind, it does so by undermining the more general thought that any purely formal process (Searle sometimes says syntactic process) could ever constitute a mind.

Clearly then, if the Chinese room argument holds, it must also hold against any 'tin-can' robot; *even one constructed in strict accord with O'Regan's extended sensorimotor principles.*

1.2 Computations and Understanding: Gödelian Arguments against Computationalism

Gödel's first incompleteness theorem states that *"any effectively generated theory capable of expressing elementary arithmetic cannot be both consistent and complete. In particular, for any consistent, effectively generated formal theory F that proves certain basic arithmetic truths, there is an arithmetical statement that is true, but not provable in the theory."* The resulting true but unprovable statement $G(g)$ is often referred to as 'the Gödel sentence' for the theory, (albeit there are infinitely many other statements in the theory that share with the Gödel sentence the property of being true but not provable from the theory).

Arguments based on Gödel's first incompleteness theorem (initially presented by John Lucas [22] were first criticised by Paul Benacerraf [3] and subsequently extended, developed and widely popularised by Roger Penrose [31] [32] [33]) typically endeavour to show that for any such formal system F, humans can find the Gödel sentence $G(g)$ whilst the computation/machine (being itself bound by F) cannot. In [32] Penrose develops a subtler reformulation of this vanilla argument that purports to show that 'the human mathematician can "see" that the Gödel Sentence is true for consistent F even though the consistent F cannot prove $G(g)$'.

A detailed discussion of Penrose's own take on the Gödelian argument is outside the scope of this introduction (for critical background see [11] and for Penrose's response see [33]). Nonetheless, it is important to note that Gödelian style arguments, purporting to show computations are not necessary for cognition, have been extensively[17] and vociferously critiqued in the literature (see [36] for a review); nevertheless interest in them - both positive and negative - continues to surface, (e.g. [45] [9]); and if such Gödelian style arguments do hold, then it is clear that they must also hold against any [virtual] O'Regan robot, *even one constructed in strict accord with sensorimotor principles.*

1.3 The 'Dancing with Pixies' Reductio

Many people hold the view that 'there is a crucial barrier between computer models of minds and real minds: the barrier of consciousness' and thus that computational simulations of mind and 'phenomenal (conscious) experiences' are conceptually distinct [48]. But is consciousness really a prerequisite for genuine cognition and the realisation of mental states? Certainly Searle believes so: "the study of the mind is the study of consciousness, in much the same sense that biology is the study of life" [41] and this observation leads him to postulate a 'connection principle' whereby, "... any mental state must be, at least in principle, capable of being brought to conscious awareness". Hence, if computational machines are not capable of enjoying consciousness, they are incapable of carrying genuine mental states and computational connectionist projects must ultimately fail as an adequate model for cognition.

In the following section I will briefly review a simple reductio ad absurdum argument that suggests there may be serious consequences in granting phenomenal (conscious) experience to any computational system purely in virtue of its execution of a particular program. If correct the 'dancing with pixies' (DwP) reductio [5] entails that either strong computational accounts of consciousness must fail OR that panpsychism is true.

The argument derives from ideas originally outlined by Hilary Putnam [37], Tim Maudlin [23], John Searle [40] and subsequently criticised by David Chalmers [12], Colin Klein [20] and Ron Chrisley [13] [14] amongst others[18].

In what follows, instead of seeking to justify Putnam's claim that "every open system implements every Finite State Automaton" (FSA) and hence that "psychological states of the brain cannot be functional states of a computer", I will simply establish the weaker result that, over a finite time window, every open physical system implements the trace of a Finite State Automata Q on fixed, specified input (I). That this result leads to panpsychism[19] is clear as, equating FSA Q(I) to a specific computational system that is claimed to instantiate phenomenal states as it executes, and following Putnam's procedure, identical computational (and, ex hypothesi, phenomenal) states can be found in every open physical system (OPS).

---

[17] For example Lucas maintains a web page http://users.ox.ac.uk/~jrlucas/Godel/referenc.html listing over fifty such criticisms.

[18] For early discussion of these themes see 'Minds and Machines': 4(4), 'What is Computation?'

[19] Panpsychism: the belief that the physical universe is composed of elements each of which is conscious.

Formally DwP is a reductio ad absurdum argument that endeavours to demonstrate that:

- IF 'an appropriately programmed computer really does instantiate genuine phenomenal states'
- THEN 'panpsychism holds'
  - *However, against the backdrop of our immense scientific knowledge of the closed physical world, and the corresponding widespread desire to explain everything ultimately in physical terms, panpsychism has come to seem an implausible view ..*
- HENCE we are led to reject the assumed claim (that an appropriately programmed computer really does instantiate genuine phenomenal states).

In his 1950 paper, 'Computing machinery and intelligence', Turing defined discrete state machines (DSMs) as "machines that move in sudden jumps or clicks from one quite definite state to another", and explained that modern digital computers fall within the class of them. An example DSM from Turing is one that cycles through three computational states (Q_1, Q_2, Q_3) at discrete clock clicks. Such a device, which cycles through a linear series of state transitions 'like clockwork', may be implemented by a simple wheel-machine that revolves through 120^0 intervals.

By labelling the three discrete positions of the wheel (A, B, C) we can map computational states of the DSM (Q_1, Q_2, Q_3) to the physical positions of the wheel (A, B, C) such that, for example, $(A \rightarrow Q_1; B \rightarrow Q_2; C \rightarrow Q_3)$. Clearly this mapping is observer relative: position A could map to Q_2 or Q_3 and, with other states appropriately reassigned, the machine's function would be unchanged. In general, we can generate the behaviour of any K-state (input-less) DSM, $f(Q) \rightarrow Q'$, by a K-state wheel-machine (e.g. a digital counter), and a function that maps each 'counter' state C_n to each computational state Q_n as required.

In addition, Turing's machine may be stopped by the application of a brake and whenever it enters a specific computational state a lamp will come on. Input to the machine is thus the state of the brake, $(I = ON|OFF)$, and its output, (Z), the state of the lamp. Hence the operation of a DSM with input is described by a series of contingent branching state transitions', which map from current state to next state $f(Q, I) \rightarrow Q'$ and define output (in the Moore form) $f(Q') \rightarrow Z$.

However, clamping the input to the device over a finite time interval entails that such input-sensitive contingent behaviour reverts to 'mere clockwork', $f(Q) \rightarrow Q'$. E.g. If Turing's DSM starts in Q_1 and the brake is OFF for two clicks, its behaviour (execution trace) is fully described by the sequence of state transitions, $(Q_1; Q_2; Q_3)$. Hence, over a finite time window, if the input to a DSM is clamped, we can map from each counter state C_n to each computational state Q_n, as required. And similarly, following Putnam, in [5] Bishop demonstrate's how to map any computational state sequence with defined input onto the [non-repeating] internal states generated by any open physical system (e.g. a rock).

Now, returning to a putative conscious robot: at the heart of such a beast there is a computational system - typically a microprocessor; memory and peripherals. With input clamped such a system is effectively a DSM with no input. Thus, with input to the robot clamped over a finite time interval, we can map its execution trace onto the state evolution of any sufficiently large digital counter or, (ibid), any OPS. Hence, if the state evolution of a robot controlled by DSM instantiates phenomenal experience, then so must the state evolution of any OPS, and we are inexorably led to a panpsychist worldview whereby disembodied phenomenal consciousnesses (aka 'pixies') are dancing everywhere ..

In [7] Bishop reviews three arguments (summarised herein) that purport to show that computations are not sufficient for cognition; for example, that the execution of a computational connectionist simulation of the brain cannot instantiate genuine understanding or phenomenal consciousness (qua computation) and hence that there are limits to the use of the computational explanations in cognitive science. But perhaps this conclusion is too strong? E.g. How do the *a priori* arguments discussed herein accommodate the important results being obtained through computational neuroscience to cognition?

There are two responses to this question. The first suggests that there may be principled reasons why it may not be possible to adequately simulate all aspects of mind via a computational system; there are bounds to a [Turing machine based] computational intelligence. Amongst others this position has been espoused by: Penrose; Copeland who claims the belief that "the action of any continuous system can be approximated by a Turing Machine to any required degree of fineness ... is false[20]"; and Smith who in [44] outlines results from 'Chaos Theory' which describe how 'Shadowing Theorems' fundamentally limit the set of chaotic functions that a Turing machine can model to those that are 'well-behaved'; i.e. functions that are not well-behaved cannot, in principle, be accurately described by Turing machine simulation.

However, Gödelian style arguments, purporting to show computations are not necessary for cognition, have been *extensively* criticised in the literature; and are currently endorsed by only a few, albeit in some cases very eminent, scholars. Nonetheless some, for example Hava Siegelmann [43], are confident that even if Gödelian arguments are valid, super Turing computation, in the form of Artificial Recurrent Neural Networks (ARNNs) offer a potential reconciliation between at least one form of [neural] computation and mind:

> Our model may also be thought of as a possible answer to Penrose's recent claim [31] that the standard model of computing is not appropriate for modelling true biological intelligence. Penrose argues that physical

[20] Copeland's argument is detailed, but at heart he follows an extremely simple line of reasoning: consider an idealised analogue computer that can add two reals (a, b) and output one if they are the same, zero otherwise. Clearly either (a) or (b) could be non-computable numbers (in the specific formal sense of non Turing-computable numbers). Hence, clearly there exists no Turing machine that, for any finite precision (k), can decide the general function $F(a = b)$, (see [15] for detailed discussion of the implications of this result).

> processes, evolving at a quantum level, may result in computations which cannot be incorporated in Church's Thesis. The analog neural network does allow for non-Turing power while keeping track of computational constraints, and thus embeds a possible answer to Penrose's challenge within the framework of classical computer science.

If Siegelmann is correct, a [O'Regan] 'sensorimotor-robot' (controlled by a suitably configured super-Turing ARNN) would be invulnerable to Penrose's Gödelian style arguments and hence, in the right context, would be capable of 'mathematical insight'.

A second response emerges from the 'Chinese room argument' and the 'dancing with pixies' *reductio*. It acknowledges the huge value that the computational metaphor plays in current cognitive science and concedes, for example, that a future computational neuroscience may be able to simulate *any* aspect of neuronal processing and offer insights into all the workings of the brain. However, although such a computational neuroscience will result in deep understanding of neuronal cognitive processes, it insists on a fundamental ontological division between the *simulation of a thing* and *the thing itself*.

For instance we may simulate the properties of gold using a computer program but such a program does not automatically confer upon us riches (unless of course the simulation becomes duplication; an identity). Hence Searle's famous observation that "*No one supposes that computer simulations of a five-alarm fire will burn the neighbourhood down or that a computer simulation of a rainstorm will leave us all drenched. Why on earth would anyone suppose that a computer simulation of understanding actually understood anything?*" [39]

Both of the above responses accommodate results from computational cognitive science, but the second specifically highlights continued shortcomings of any **purely formal** - vis a vis computational - account of cognition, such as, we suggest, Kevin O'Regan's conception of sensorimotor theory (contra Alva Noë's) is committed to. If this analysis is correct, then it is perhaps time for contemporary sensorimotor theorists to embrace the 'strong-embodiment of brain and body' (and concomitant physical and social context) a little more deeply[21] ..

2 A Brief Resume of 'Contemporary Sensorimotor Theory'

We believe this survey of contemporary sensorimotor theory offers an interesting selection of current research informed by the sensorimotor account of perception and we will complete this introduction with a brief summary of the contributed works. In this conception we are privileged to open our volume with an introduction to contemporary sensorimotor theory from J. Kevin O'Regan.

In his opening chapter Kevin offers a new overview of extant sensorimotor theory with particular regard to the hard problem of 'phenomenological consciousness'. In so doing, he demonstrates the power of a method - fundamentally

[21] C.f. Gibson [18], Varela et al [49], Bickhard & Terveen [4], Thompson [46], Deacon [16] etc.

grounded in observable physical phenomena - of describing and potentially explaining features of consciousness that have often been regarded as falling outside of the scope of genuine scientific explanation.

After this resume of contemporary sensorimotor theory, the field is critically scrutinised from two distinct perspectives: (a) 'the body, enactivism and emotion' and (b) 'action as constitutive of perception'.

Firstly Paine describes (a) how sensorimotor theory can be shown to have 'Heideggerian roots' (such as its prioritisation of environmental engagement) and (b) the degrees to which it satisfies a range of Heideggerian conditions for the presence of conscious experience; in the process Paine explores key differences between the two approaches and offers suggestions as to how they might be resolved. Paine concludes that although O'Regan's 2011 sensorimotor account [29] meets Dreyfus' requirements for a 'Heideggerian AI' (and in the process potentially satisfies his objections to 'good old fashioned - classical - Artificial Intelligence'), it 'omits room for emotion'; an area where, she suggests, a Heideggerian perspective would offer further insight.

Subsequently, and in a very carefully argued piece, Ken Pepper draws on the philosophy of Maurice Merleau-Ponty to sketch a 'phenomenological interpretation' of sensorimotor understanding, appealing to the 'Phenomenology of Perception' to show (a) how two of its major operative concepts - the 'body schema' and 'sedimentation' - can help plug the gaps in Noë's sensorimotor account and (b) how the notion of 'body schema' conditions the way things appear to the agent. In so doing Pepper presents a detailed analysis of the differences between 'affordances' and 'object horizons', and clarifies Noë's position with respect to the division.

Then Scarinzi looks at how enactive - with respect to the role of mental representations and embodiment - the sensorimotor approach really is. Scarinzi argues that it is actually 'semi-enactive' and that to bring it closer to enactivism proper (in investigation of the agent's subjective experience of the qualities of interactions) it needs to be more sensitive to (i) the motor and cognitive-emotional role of the lived body and (ii) the agent's subjective access to it. Thus, although sensorimotor theory and enactivism proper share some important features (such as a fundamental entanglement with the environment) sensorimotor theory merely skirts other experiential features (such as the role of emotional involvement); in this way Scarinzi suggests that in the sensorimotor account *the subjectivity of lived experience* is neglected.

Secondly Rainey (and subsequently Loughlin) examine the so called 'causal - constitutive objection':

> Noë argues that visual states are not pictorial; he argues that all perception is conceptual; and he argues that the external world makes a constitutive contribution to experience. I am unpersuaded by these arguments .. (Jesse Prinz [35]).
>
> Even if perceptual experience depends causally or counterfactually on movement or another form of activity, it does not follow that perceptual experience constitutively involves movement .. (Ned Block [8]).

Thus Rainey draws on various philosophical traditions to consider the controversial and potentially problematic claim - in the context of sensorimotor theory - that experience is conceptual; i.e. that experience is a complex, conceptually articulated and conceptually very laden affair. If Rainey is correct, and sensorimotor theory is indeed committed to a conceptual view of experience, then he argues, it opens a Pandora's box for critics like Block and Prinz to argue the *causal-constitutive objection.*

Loughlin suggests that sensorimotor theory is 'centrally committed to the claim that visual experience is realized by embodied know-how or skilful engagement' which, Loughlin claims, opens it up to attacks from Aizawa, Block, Clark, Prinz et al. That said, Loughlin concludes that sensorimotor theory can easily accommodate most of their objections, but that the most serious - the causal-constitutive objection - can only be avoided by 'going radical', i.e. shifting to a conceptual stance more closely aligned with Hutto and Myin's 'radical enactivism' [19].

Further scrutiny of the sensorimotor account takes the form of identifying and describing problems in the theory as it stands, and identifying related approaches by which these challenges could potentially be met. In this vein, Wadham identifies an issue in Noë's development of sensorimotor theory, claiming his account of perceptual content as being 'virtual all the way in' is incompatible with his account of perspectival content via 'p-properties'. Wadham suggests that the issue arises from the 'problem of invisible contents'; i.e. Noë's virtual content thesis implies that p-properties must be invisible[22] whereas the key role these properties play in Noë's overarching theory of perception requires that they cannot be 'invisible' in this way[23]. Wadham concludes by offering a potential resolution to this contradiction, suggesting an alternative to Noë's p-property story: 'appearance-pattern theory'.

Subsequently Lyon suggests that classical sensorimotor accounts of perception, based mainly on vision and touch, are inadequate for other sensory modalities - specifically audition. Lyon examines the effect of including sound in accounts of perception, and suggests that it makes sense to avoid the 'unnecessary strait jacket' of a model based primarily on vision and touch alone. Lyon concludes by suggesting that the sensorimotor approach can be usefully extended to other perceptual modes.

Parthemore considers how sensorimotor theory aligns with and differs from the unified conceptual space theory of concepts (as derived from Gardenfors' canonical conceptual space theory' [17]), showing how the meta-theory resulting from a combination of both introduces a new role for emotional affect and

[22] "*.. while in motion, we cannot see individual p-properties .. if we can't take in any content in an instant, then we can't see a property of the object that is presented to us only for an instant"; hence p-properties are invisible*", (Wadham, this volume).

[23] "*.. as we move, we experience a variety of p-properties presented by the object we are looking at. It is through implicitly understanding the way in which the p-properties we see vary as a function of our movement that we come to see the actual, perspective-independent properties of objects*", (Wadham, this volume).

the somatosensory system. These additional aspects are suggested as offering a potential avenue for grounding both 'salience' and the 'proper understanding of mental representation'.

In an echo of Tononi [47] and Aleksander & Morton [1], Gamez examines sensorimotor theory through an information-theoretic lens that views the character of sensation as being constituted by the way in which 'information is processed and structured'. In this light he offers critical perspectives on three extant issues sometimes viewed as problems for the sensorimotor account: (a) the view that sensation may exist without motor action; (b) conscious sensations can be evoked by directly stimulating the brain, and (c) the claim that visual sensory substitution fails to produce *genuine* visual sensations. Then, in controversial opposition to central tenets of sensorimotor theory, Gamez endorses a weaker form of theory that merely posits a *correlation* between sensory or sensorimotor dependencies and the contents of consciousness[24] - effectively "bridging" between classical (correlational) accounts of consciousness and strict sensorimotor accounts - such that (a) conscious sensations are merely correlated with sensory dependencies, and (b) conscious perception is merely correlated with mastery of sensorimotor dependencies.

In the last section of the book we examine how contemporary sensorimotor theory has been successfully applied in other domains. Thus, the key claim of Rucinska's paper is to develop a theory of basic play grounded in 'action' rather than 'representations'. In so doing Rucinska makes a convincing case that sensorimotor affordances can be seen as a sufficient basis for imaginative play in young children; specifically, Rucinska considers how a sensorimotor account of perception can be used to form a 'theory of basic pretence' considering, for example, the possibility that a child may pretend that a banana is a phone through direct on-line capacities, rather than requiring off-line conceptual workings, and furthermore without requiring internal representations. As O'Regan and Noë described 'perceiving as a way of acting', so Rucinska suggests 'pretending as a way of acting'. In shifting focus to applying sensorimotor theory as it stands to a new area of investigation, its usefulness is demonstrated as a force both directing research and informing explanation of results.

Gibbs and Devlin present a survey of sensory substitution and sensory augmentation devices, and enactive interfaces in real, as well as virtual, worlds. They show how a sensorimotor interpretation of existing work on sensory augmentation can be applied to make predictions about the experience of reproducing the experiments in virtual worlds. These predictions, being testable, provide a basis for research into both perception and the explanatory power of the sensorimotor account itself.

Similarly Gillies & Kleinsmith investigate how the sensorimotor account of perception can inform user-interface design, suggesting that learned interactions

[24] "*While some authors have suggested that there is an identity between a mastery of sensorimotor contingencies and consciousness .. I am only focusing here on the weaker and less contentious claim that there might be a correlation between sensory or sensorimotor contingencies and the contents of consciousness*", Gamez (herein).

can be exploited to allow someone to design with performed actions, rather than representations thereof. Gillies & Kleinsmith illustrate this conception via an innovative application used to design interactions with video game characters.

Hoffmann discusses a number of case studies of robots that detect and exploit sensorimotor regularities in their environment to exhibit more complex behaviour than has previously been achievable. He subsequently broadens his perspective to draw out more general aspects of the relation between body schema, forward models and sensorimotor contingency theory, concluding by highlighting potential limitations in the field of cognitive robotics.

Finally Cowley asks if current conceptual shortcomings in dealing with language result primarily from a fixation with the written word, and recommends a more enactive approach to the subject, rather than, for example, a continuing focus on Chomsky's mentalism. Nonetheless, Cowley warns against a simple reduction to 'sensorimotor dependencies' as explanatory of linguistic behaviour: "*.. in embodied cognition, there is a risk of overplaying work on how bodies (and brains) regulate activity/ system-states. In ORegan and No [27], for example, perceptual modalities are said to exist 'only in the context of the interacting organism'*". A case study - emphasising the dual role of action and language in perception - is examined in detail; it clearly illustrates that such a simple reduction does not apply in human languaging[25].

As editors we think this collection of new essays offers a fascinating snapshot of contemporary sensorimotor theory in 2014; early leaves from O'Regan and Noë's 'enactivist approach' have blossomed ..

Acknowledgements. JMB would like to thank his wife Dr. Katerina Koutsantoni for her enduring patience in preparing this volume, his co-editor, Andrew Martin, who has worked miracles in helping to bring this volume about and the reviewers whose careful comments helped both polish the content of the volume and clarify several important issues. He would also like to specifically thank Jan Degenaar, Kevin O'Regan and Alva Noë for their very helpful comments in preparing this introduction.

AOM would like to thank his parents and Norah who contributed much of their time to support the work that went into this volume and the unerring guidance of his co-editor, supervisor and friend 'Bish', who conceived of, and initiated, the project.

Many of the papers herein were originally presented at the "1st AISB Workshop on Sensorimotor Theory" held at Goldsmiths, University of London, on September 26th, 2012, and thus thanks must also be extended to the AISB[26],

[25] "*.. while the young mans thinking is based in sensorimotor activity, his actions reach into a social domain that lies beyond body-world interaction*, Cowley (herein)."

[26] The AISB is the UK society for the study of Artificial Intelligence and the Simulation of Behaviour. It is the longest established such society in the world and celebrates its 50th anniversary at a Convention at Goldsmiths in April 2014.

EUCOG[27] and Goldsmiths, University of London for their generous support of that event. In preparing this introduction we have extracted the description of CRA from Bishop [6] and DwP from Bishop [7].

References

1. Aleksander, I., Morton, H.: Aristotles laptop: the discovery of our informational mind. World Scientific Publishing, Singapore (2012)
2. Andersson, R.L.: Robot ping-pong player. The MIT Press, Cambridge (1988)
3. Benacerraf, P.: God, the Devil & Gödel, Monist 51 (1967)
4. Bickhard, M.H., Terveen, L.: Foundational Issues in Artificial Intelligence and Cognitive Science. North Holland (1995)
5. Bishop, J.M.: Dancing with Pixies: strong artificial intelligence and panpyschism. In: Preston, J., Bishop, J.M. (eds.) Views into the Chinese Room: New Essays on Searle and Artificial Intelligence. Oxford University Press, Oxford (2002)
6. Bishop, J.M.: A view inside the Chinese room. Philosopher 28(4), 47–51 (2004)
7. Bishop, J.M.: A Cognitive Computation fallacy? Cognition, computations and panpsychism. Cognitive Computation 1(3), 221–233 (2009)
8. Block, N.: Review of Alva Noë, Action in Perception. Journal of Philosophy 102, 259–272 (2005)
9. Bringsjord, S., Xiao, H.: A refutation of Penrose' Gödelian case against artificial intelligence. J. Exp. Theo. Art. Int. 12, 307–329 (2000)
10. Carruthers, P.: Higher-Order Theories of Consciousness. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy (2011)
11. Chalmers, D.: Minds, Machines, And Mathematics A Review of Shadows of the Mind by Roger Penrose. PSYCHE 2(9) (1995)
12. Chalmers, D.J.: The Conscious Mind: In Search of a Fundamental Theory. Oxford University Press, Oxford (1996)
13. Chrisley, R.: Why Everything Doesn't Realize Every Computation. Minds and Machines 4, 403–420 (1995)
14. Chrisley, R.: Counterfactual computational vehicles of consciousness. In: Toward a Science of Consciousness, Tucson Arizona, USA, April 4-8 (2006)
15. Copeland, B.J.: The broad conception of computation. American Behavioral Scientist 40(6), 690–716 (1997)
16. Deacon, T.: Incomplete Nature: How Mind Emerged from Matter. W. W. Norton & Company (2012)
17. Gardenfors, P.: Conceptual Spaces as a Framework for Knowledge Representation. Mind and Matter 2(2), 9–27 (2004)
18. Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston (1979)
19. Hutto, D., Myin, D.: Radicalizing Enactivism: basic minds without content. The MIT Press, Cambridge (2012)

[27] EUCog is a European network of nearly 900 researchers in artificial cognitive systems and related areas who want to connect to other researchers, reflect on the challenges of the discipline and get their research 'out there'. The network funds meetings, workshops, members' participation in academic events, faculty exchanges and other activities that further its aims. It is funded by the Information and Communication Technologies division of the European Commission, Cognitive Systems and Robotics unit, under the 7th Research Framework Programme.

20. Klein, C.: Maudlin on Computation, working paper (2004)
21. Levine, J.: Materialism and qualia: the explanatory gap. Pacific Philosophical Quarterly 64, 354–361 (1983)
22. Lucas, J.R.: Minds, Machines & Gödel. Philosophy 36 (1961)
23. Maudlin, T.: Computation and Consciousness. Journal of Philosophy 86, 407–432 (1989)
24. Nagel, T.: What Is It Like to Be a Bat? The Philosophical Review 83(4), 435–450 (1974)
25. Noë, A.: Action in Perception. The MIT Press, Cambridge (2004)
26. Noë, A.: Out of Our Heads: why you are not your brain, and other lessons from the biology of consciousness. Hill and Wang, NYC (2009)
27. O'Regan, K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behaviour and Brain Sciences 24, 939–1031 (2001)
28. O'Regan, J.K.: Sensorimotor approach to (phenomenal) consciousness. In: Baynes, T., Cleeremans, A., Wilken, P. (eds.) Oxford Companion to Consciousness, pp. 588–593. Oxford University Press, Oxford (2009)
29. O'Regan, J.K.: Why red doesn't sound like a bell: understanding the feel of consciousness. Oxford University Press, Oxford (2011)
30. O'Regan, J.K.: How to Build a Robot that is Conscious and Feels. Minds and Machines 22(2), 117–136 (2012)
31. Penrose, R.: The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics. Oxford University Press, Oxford (1989)
32. Penrose, R.: Shadows of the Mind: A Search for the Missing Science of Consciousness. Oxford University Press, Oxford
33. Penrose, R.: Beyond the Doubting of a Shadow A Reply to Commentaries on Shadows of the Mind. PSYCHE 2(23) (1996)
34. Preston, J., Bishop, J.M. (eds.): Views into the Chinese Room: New Essays on Searle and Artificial Intelligence. Oxford University Press, Oxford (2002)
35. Prinz, J.: Putting the Brakes on Enactive Perception. PSYCHE 12(1) (2006)
36. Psyche: Symposium on Roger Penrose's Shadows of the Mind. PSYCHE 2 (1995)
37. Putnam, H.: Representation & Reality. Bradford Books, Cambridge (1988)
38. Rosenthal, D.M.: Varieties of higher-order theory. In: Gennaro, R.J. (ed.) Higher-Order Theories of Consciousness: An Anthology. Johns Benjamins Publishing Company, Amsterdam (2004)
39. Searle, J.R.: Minds, brains, and programs. Behavioral and Brain Sciences 3(3), 417–457 (1980)
40. Searle, J.: Is the Brain a Digital Computer? Proceedings of the American Philosophical Association 64, 21–37 (1990)
41. Searle, J.: The Rediscovery of the Mind, p. 227. MIT Press, Cambridge (1992)
42. Searle, J.: The Mystery of Consciousness. Granta Books, London (1994)
43. Siegelmann, H.T.: Neural and Super-Turing Computing. Minds and Machines 13, 103–114 (1993)
44. Smith, P.: Explaining Chaos. Cambridge University Press, Cambridge (1998)
45. Tassinari, R.P., D'Ottaviano, I.M.L.: Cogito ergo sum non machina! About Gödel's first incompleteness theorem and Turing machines. CLE e-Prints 7(3) (2007)
46. Thompson, E.: Mind in Life: biology, phenomenology, and the sciences of mind. Harvard University Press, Cambridge (2010)
47. Tononi, G.: Consciousness as integrated information: a provisional manifesto. Biol. Bull. 215, 216–242 (2008)

48. Torrance, S.: Thin Phenomenality and Machine Consciousness. In: Chrisley, R., Clowes, R., Torrance, S. (eds.) Proc. 2005 Symposium on Next Generation Approaches to Machine Consciousness: Imagination, Development, Intersubjectivity and Embodiment, AISB05 Convention, pp. 59–66. University of Hertfordshire, Hertfordshire (2005)
49. Varela, F.J., Thompson, E., Roesch, E.: The Embodied Mind: Cognitive Science and Human Experience. MIT Press, Cambridge (1991)
50. Warwick, K.: Alien Encounters. In: Preston, J., Bishop, J.M. (eds.) Views into the Chinese Room: New Essays on Searle and Artificial Intelligence. Oxford University Press, Oxford (2002)
51. Wheeler, M.: Change in the Rules: Computers, Dynamical Systems, and Searle. In: Preston, J., Bishop, J.M. (eds.) Views into the Chinese Room: New Essays on Searle and Artificial Intelligence. Oxford University Press, Oxford (2002)

The Explanatory Status of the Sensorimotor Approach to Phenomenal Consciousness, and Its Appeal to Cognition

J. Kevin O'Regan

Laboratoire Psychologie de la Perception - Université Paris Descartes, Paris, France
jkevin.oregan@gmail.com
http://www.kevin-oregan.net

Abstract. This paper starts by providing a succinct overview of the sensorimotor approach to phenomenal consciousness, describing its two parts: the part that concerns the quality of sensations, and the part that concerns whether or not such qualities are (consciously) experienced. The paper goes on to discuss the explanatory status of the approach, claiming that the approach does not simply "explain away" qualia, but that on the contrary, it provides a way of thinking about qualia that explains why they are the way they are, stimulates scientific paradigms and produces testable predictions. A final part of the paper examines the relation of the theory to radical enactivism, claiming that some kind of "higher order" cognitive mechanism similar to that used in Higher Order Thought theories of consciousness is needed to account for what is usually meant by being conscious of something.

Keywords: Qualia, phenomenal consciousness, radical enactivism, Higher Order Thought, sensorimotor theory.

1 Introduction

In this paper I will address two contentious points about the sensorimotor approach to consciousness. These concern the explanatory status of the approach, and the role and nature of cognition in the approach. But first a quick summary of the main ideas in the approach.

2 Quick Summary of the Sensorimotor Approach

The "hard" problem of phenomenal consciousness is the problem of *qualia*, that is, of understanding what it might be about, for example, red-sensitive circuits in the brain that causes them to create a "red" feel rather than, say, a "green" or "onion flavor" feel, or even any feel at all. Another way of illustrating the problem is to ask: "What would we have to build into a robot so that it really felt the touch of a finger, the redness of red, or the hurt of a pain?" The sensorimotor approach ([1], [2]) answers these questions by thinking about experience in a

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_2,

new way. The approach is not just a philosophical theory but has scientific value in that it has been successful in driving new research paradigms.

The approach starts by making a distinction, and invoking separate mechanisms, to explain (1) the experienced quality of a feel, and (2) whether or not a person is conscious of this quality.

2.1 The Experienced Quality

Instead of thinking of the experienced quality of a sensory experience as being generated somewhere in the brain, the sensorimotor approach contends that the quality is constituted by the set of objective laws concerning the interaction with the world that the experience involves. The objective laws linking actions to resulting sensory changes are called "sensorimotor contingencies" or "sensorimotor dependencies".

For example, the quality of tactile experience is determined by the laws that determine the changes of sensory input that occur when you move your body with respect to a stimulating surface: the softness of a sponge, for example, is constituted by the fact that when you press on the sponge it cedes under your pressure. The quality of auditory experience is determined by laws like the fact that when you approach a sound source, the amplitude of the sensory input increases, etc. The similarities and differences between experienced qualities are constituted by the similarities and differences in such sensorimotor dependencies.

In addition to explaining similarities and differences between qualities, the view also explains "what it's like" (in the expression of Nagel [3]) to have a *sensory* experience: Certain objective physical facts about real sensory interactions are unique to the classic five sensory systems of vision, hearing, touching, tasting and smelling, and do not apply to mental activities like thoughts or imaginings, nor to autonomic processes in the nervous system. These objective facts are bodiliness, insubordinateness and grabbiness. Bodiliness is the objective fact that any movement of the body immediately changes input coming from sensory receptors. Insubordinateness is the objective fact that sensory input can be changed by the outside world without voluntary control by the person. Grabbiness is the objective fact that sudden changes in sensory input channels automatically capture attention. Bodiliness, insubordinateness and grabbiness provide sensory experiences with a quality of imposing themselves on us, of partly escaping our voluntary control. Taken together, they can be put into correspondence with the notion of "reality" or sensory "presence" which is typical of the phenomenal quality of sensory experiences. Bodiliness, insubordinateness and grabbiness, when plotted on a "phenomenality plot" predict the extent to which people experience sensations as possessing this quality of having "something it's like" ([1], [4]).

2.2 Consciousness of the Quality

While sensorimotor contingencies determine the *quality* of an experience, including its sensory presence, a second ingredient is needed for that quality to be perceived *consciously*. Here the sensorimotor approach appeals to a particular, hierarchical, form of cognitive access similar to that used in the "higher-order" theories (HOT) of consciousness ([5], [6]). Returning to the example of the softness of a sponge, a person who was absent-mindedly washing the dishes with a sponge while they were talking to a friend would not at that moment be consciously experiencing the softness. Only if they cast their attention on and cognitively accessed the quality of their interaction with the sponge would they normally be said to be consciously experiencing the softness. I shall come back to this cognitive access later.

3 Solving or Dissolving the Hard Problem?

The trick used in the sensorimotor approach is to try to "dissolve" the hard problem of qualia rather than "solving" it. The idea is that the so-called hard problem is only hard because the terms generally used to talk about consciousness are unclear. When consciousness is given a clear definition, the problem dissolves. In particular, as regards phenomenal experience, the sensorimotor approach suggests that if we think clearly about what we mean by having an experience, then we see that it involves, not something being generated in the brain, but *a way of interacting with the world.* The hard problem of finding a brain mechanism that generates experience is thereby eluded.

Note however that it is not the case that the sensorimotor approach "explains away" qualia, in the sense that it brushes them aside as a non-problem. It takes seriously the question of accounting for the experienced quality of sensation, and provides answers that are scientifically testable.

The case of softness is a case where this approach works well. Consider for example the "hard" question of how the softness of the sponge might be generated by brain mechanisms. From the point of view of the sensorimotor theory, this question does not make sense, because the softness *lies in the fact* that if I press the sponge, it squishes under my pressure. It is a sensorimotor law. As an abstraction, it cannot be found anywhere in particular, and certainly it is not "generated" in the brain. (Of course the brain has a role to play, namely to enable all the mechanisms that allow the organism to be engaged with the sensorimotor laws of softness – but there will be many such mechanisms, and they will probably be widely distributed in the brain, since all sorts of motor control and sensory systems underlie our ability to interact in the "soft"-specific way. And no one, or combination of, these mechanisms could be considered to "generate" the softness feel.)

If we ask the "hard" question of why softness feels the way it does, the sensorimotor approach says that the answer boils down to the objective physical facts that describe the interaction we have with the sponge. The possibility of inverted qualities – the idea that somehow a soft sponge would feel hard and a hard sponge would feel soft – becomes inconceivable because what we *mean* by feeling softness by definition requires the sponge to react appropriately to our pressing, and this can only occur if objectively the sponge is soft.

And finally we could ask why softness has sensory presence at all, rather than providing an experience like, say, a thought or like imagining. For this the sensorimotor approach appeals to the concepts of bodiliness, insubordinateness and grabbiness. These supply potentially available facts about the current ongoing interaction which are constitutive of its "realness" or "presence". Because they have the consequence that our minds are subservient to the external stimulation, the quality of softness imposes itself upon us and manifests its presence in a way that other mental processes like thoughts, imaginings, memory or autonomic processes do not.

Thus, in the case of the softness of the sponge, important aspects of the explanatory gap dissolve, and we have an intuitively plausible way of accounting for not only the quality itself, but also its sensory presence.

However a priori it does not seem obvious that all sensations can be accommodated under this view. In particular, sensations like those of colour and smell, and perhaps even auditory and passive tactile sensations do not seem to involve actively interacting with the world. It is nevertheless the wager of the sensorimotor theory to try to cover *all* sensations in this way. We do this because of the philosophical advantage the approach promises in bridging the explanatory gap.

And furthermore, when the idea is taken seriously, it leads to productive science: stimulating the discovery of change blindness, accounting for phenomena in color science, and promoting research in sensory substitution. In a way it is surprising that finding proper definitions should propulse new scientific advances. The reason is presumably that clear definitions allow the relations between objective phenomena to be more clearly seen, revealing causes and effects that are otherwise obscured.

4 The Explanatory Status of the Approach, and Anthropocentric Fears

What then is the explanatory status of the sensorimotor approach? One should probably not say that the sensorimotor approach is a *theory* in the sense that it provides explanations. Rather, it is a way of thinking about the phenomena of consciousness that allows them to be divided into scientifically amenable bits, and which preserves a plausible link with some of the important ways we use the term in everyday life. Questions that seemed to demand explanations evaporate when we think in the new way.

By way of analogy, one could say that the sensorimotor account answers the question of phenomenal consciousness in a similar way to how the question of life was answered in modern biology. Instead of finding a magical vital spirit that generated life, modern biology discovered that it was a "category mistake" (cf. Ryle [7]) to look for a substance to generate life. We now conceive of life as being an *abstraction used to describe organisms that interact in certain ways with their environments.* By taking this stance, modern biology now decomposes the problem of life into the smaller problems of accounting for each of the aspects of the way organisms interact with their environments. And each of these subproblems becomes tractable using the normal laws of physics and chemistry. No magic vital essence is required. Of course poets will say that the magic of life is not fully captured.

Similarly the sensorimotor approach considers that it is a category mistake to search for a neural correlate that generates consciousness in the brain. Instead, the sensorimotor approach suggests that it is advantageous to more carefully define the different aspects of consciousness, and decompose it in two parts: the sensorimotor interactions whose laws constitute experienced quality; and a form of cognitive access which makes that quality conscious. Both parts are amenable to a biological explanation, and both parts keep a link with normal everyday notions of consciousness, with some limitations (see the section on ethical consequences).

The trouble is that just as it was difficult for many people to accept that modern biology could adequately account for the “mystery of life”, many people resist the sensorimotor approach, and claim that it cannot dissolve the “mystery of qualia”.

I suggest that in both cases however the problem lies not in the theories, but in an anthropocentric fear. At the beginning of the 20th Century people felt threatened by the idea that “life” might just be a chemical process. Today, people are threatened by the idea that “feel” is simply a way of interacting with the world. After all, humans are not the only agents that can interact with the world, and even machines could be said to do so. There is the terrible danger that feel, this very special property of sentient beings, might be usurped by machines. Subconsciously it is therefore an excellent tactic to entertain a degree of artistic confusion about what consciousness is, and to resist pidgeonholing it in any rational way, and in particular in the way suggested by the sensorimotor approach. That way we reduce the chances that we will ever be labelled as machines. Thus many scientists, and among them even roboticists and workers in artificial intelligence, believe that there is something fundamentally missing in our understanding of human feel, and which is preventing us building it into our machines today. I claim that this belief constitutes an unconsciously motivated diversion tactic to preserve our last bastion of humanity, namely feel, from machines. It is a form of anthropocentric chauvinism, a hidden fear that secretly fuels an unconsciously maintained confusion.

5 The Need for Cognitive Access

If we accept the wager that sensorimotor laws will completely describe all aspects of the phenomenal quality of our experiences, then another issue arises. Does the theory really need the second mechanism, namely the mechanism of cognitive access? Hutto & Myin ([8]) consider that invoking cognitive access is a capitulation to intellectualist, cognitivist and representationalist theories. They recommend dispensing with cognitive access and propose taking a radical enactivist position in an account of consciousness that retains only the sensorimotor contingencies. More representation-friendly theorists, such as Block (cf. [9]) have further intimated that in any case the appeal to cognitive access does not solve the explanatory gap problem, and sneaks in the concept of consciousness by the back door.

5.1 Radical Enactivism

There would appear to be two reasons for which a radical enactivist might want to deny the role of cognitive access in determining phenomenal consciousness. The first might be to claim that all cognition is ultimately based in sensorimotor interaction with the world. Thus the higher order form of cognitive access that the sensorimotor theory introduces to characterise consciousness of sensory qualities should be considered to be "enactive", non-representational processes. As we shall see in the next section, I have no objection to such a view, provided it retains the notion that there is the enactive equivalent of "higher order" access that determines whether the quality of the interaction is experienced as conscious. Further, perhaps the higher and lower order access should not be considered as a dichotomy, but as two extremes in a continuum of access.

A second reason to deny the role of cognitive access in determining phenomenal consciousness might be essentially a matter of definition. It might be argued that there are forms of phenomenal consciousness which do not require the perceiver to be cognitively occupied with ongoing sensorimotor activity. Rather, this sensorimotor activity could be considered to provide phenomenal consciousness in itself. Proponents of such a view might be motivated by the intuition that even when I'm talking to a friend while I wash the dishes, I am still in some sense experiencing the softness of the sponge. Further, presumably beings like newborn babies or animals with limited cognitive capacities – let's say frogs – do really experience their worlds phenomenally. Even without well-developed minds, such beings nevertheless possess some kind of minimal, "organismic" form of feel ([10], [11]).

Taking this option of defining certain kinds of sensorimotor interactions to possess phenomenal consciousness in themselves, leads down a slippery slope. Consider a system like our digestive system, the immune system, or plants. One could also look at artificial systems like missile guidance systems or even the lowly thermostat. All these systems could be considered sensorimotor systems. They engage in sensorimotor interactions. One could say that they exercise sensorimotor laws. But surely they do not feel anything.

One could argue that such systems would begin to have experiences if one made them sufficiently complex. Perhaps the nematode worm c. elegans, which is probably a kind of biological automaton with a fixed number of cells and well-defined neural pathways has no feel, but a fly, or if not, perhaps a frog, or a mouse surely has some kind of feel. Intuitively this might be because it is not completely pre-wired, because it has a certain degree of neural plasticity and can learn from its past experiences, or because it is sufficiently complex. But why? What might it be about not being pre-wired, about plasticity or complexity that instills feel into a system? Or you could argue that it is simply what you *mean* by saying that a system has a feel, that the system displays a certain amount of adaptability, flexibility and complexity. We could take this stance *as a matter of definition* for feel, but then all sorts of systems, be they biological, mechanical, economical, evolutionary, display flexibility, adaptability and complexity, and yet presumably they don't feel. Perhaps certain biological systems (namely ourselves) have an extra special degree of flexibility, adaptability and complexity. But it is not clear what exactly this extra special degree is, nor how to characterise it in a way that makes it evidently constitutive of feel.

Finally another tactic might be to conjecture that in order to provide such special adaptability, flexibility and complexity, some kind of special brain systems are necessary, and these produce feel *as a by-product*. An obvious candidate would be systems involved in drives, emotions or arousal. But such arguments are precisely the arguments that the sensorimotor theory was trying to avoid. They hark back to the "category mistake" of thinking that feel is a substance produced by the brain, and they immediately face the explanatory gap problem: how could one *logically* ever explain how such emotional or arousal systems might *create* feel. Whatever chemical or neural mechanism was postulated, be it some neuromodulator in the brain that labels states as being emotionally charged, or be it the likes of oscillations, neural synchronisation, re-entrant loops, or even quantum gravity mechanisms in microtubules, one would have to explain precisely *what it is* about this mechanism that made it "like something" to feel.

The solution proposed by the sensorimotor approach is therefore to reject the existence of an "organismic" kind of feel that requires only non-cognitive sensorimotor interactions, and to claim that sensorimotor interactions are only part of what we mean by experiencing. Certainly their qualities constitute the experienced phenomenal qualities. But a second mechanism, namely a form of higher order cognitive access, is additionally required to make the experienced qualities *conscious*. In fact as we shall see below, it could be argued that what I call higher order cognitive access is actually just a way of being a little more precise about the notions of adaptability, flexibility and complexity which intuitively we want to have in our definition of consciousness. In any case I maintain that something resembling cognition, as well as the idea of different modes of access to ongoing interactions, are necessary to capture a satisfactory definition of what we generally mean by consciousness.

5.2 Adopting a Form of Higher Order Thought

To support the idea that some notion of higher order cognition is essentially *what people normally mean* when they say they are conscious of something, let us consider again the situation where I am washing the dishes with a sponge. Let us try to capture what it *means* to say that I'm conscious of the softness of the sponge.

What first comes to mind is to say that:

(A): In my current rational behaviour (such as deciding or planning actions, possibly making linguistic utterances), I am making use of the fact that I can currently confirm that the sponge is soft.

We shall see that this is not sufficient and needs further development, but let us use this as a working start. And let us decompose each of the underlying statements, starting at the end.

Saying that I can currently confirm that the sponge is soft means that I can confirm that when I press it, it squishes. But there are many ways of pressing, and many ways of confirming that it squishes, involving different muscle combinations, finger positions, etc., each involving different neural signals. The details are not relevant however, because my brain has abstracted a category called "soft", and the current sensorimotor activity fits this category. Furthermore when I say: "I can confirm...", this does not mean that I as a person am aware of how this abstraction has been effected, nor does it mean that I am aware of the intimate neural, muscular or sensory workings involved. On the contrary, it is the case that my brain must have made this abstraction, and must now be confirming that its conditions are fulfilled. In past papers, I have described this situation by saying that the brain must "currently be exercising its mastery of the sensorimotor contingency of softness".

So to be more precise, we should modify (A) as follows:

(A'): In my current rational behaviour (such as deciding or planning actions, possibly making linguistic utterances), I am making use of the fact that my brain is currently exercising mastery of the sensorimotor contingency of softness.

Note also that because softness is an abstract sensorimotor law, statement (A') says that I am making use of "the fact" that I can confirm that the sponge is soft: I'm not making use of a set of neural commands, or an ongoing pattern of neural activities or muscular give-and-takes: I am making use of *the fact* that a condition has been satisfied.

Now let us analyse more closely what it means to be "making use, in my rational behaviour" of this fact. It means that the fact-of-softness is playing a role in this behaviour. For example, I have grasped the sponge in order to wipe the plate, rather than trying to use the plate to wipe the sponge. Dishwashing is a rational behaviour, requiring, among other things, to be able to use the softness to appropriately direct sponge-use.

But notice that I could be doing the dishwashing as I'm absorbed in following an interesting TV programme. Although it would be true to say that I'm making rational use of the sponge's softness, under normal parlance one would not say I was conscious of the softness of the sponge. Furthermore one could imagine a dishwashing robot that appropriately chose sponges to wash dishes (rather than the other way around), doing so by detecting that the sponges are soft whereas dishes are hard. Statement (A') would then apply to the robot, but in normal parlance we would not say that the robot was conscious of the softness.

For us to be more willing to admit that the robot was conscious of the softness, it would have to be not only making use of the softness in its dishwashing, but, as suggested in HOT theories ([5], [6]), it would have to *know* it was doing this. This kind of knowing would have to be of a fairly high level: thus it would not be sufficient if the robot just wrote out on its terminal "Currently checking softness to choose between sponge and plate...", since this kind of statement could be generated automatically by its dishwashing routine. To appear conscious, the robot would have to give us evidence that it knew what it was doing *in a wider context*. It should "potentially have in mind" that there are a variety of other things that it might have been mentally accessing instead of the softness: and these things would not only be those that are part of the dish-washing routine like the hotness of the water, the size of the plates, etc.; they would include contextual things like the sound of the television in the background, the tiredness in its limbs; and even more general thoughts or attitudes the robot might have towards what it is doing and its current situation, like the fact that today there are more dishes than usual, or that it might be more pleasant to be reading a book, etc. I suggest that in normal parlance, being conscious of the softness of the sponge involves not just attending to the softness of the sponge, but doing so in a way that includes knowing that potentially one could also attend to a variety of related things that are "latent" or in the "fringe" of one's awareness.

Let us call this "wider" type of "making use in one's rational behaviour of X": "having higher order cognitive access to X". This notion of higher order cognitive access thus implicitly includes not just X, which is the central thing actually being "made use" of, but also a range of other "latent" things in the "fringe", say W,Y,Z..., and that pertain to one's current situation, to the current context, or to one's thoughts or attitudes towards X, and which can *potentially* also be "made use" of.

Then, a better approximation than (A') to *what I mean* by being conscious of the softness of the sponge is:

(B) I have higher order cognitive access to the fact that my brain is currently exercising mastery of the sensorimotor contingency of softness.

In [1] I had formulated (B) by employing the term "cognitive access" in a hierarchical way: I had said that what I mean by being conscious of the softness of the sponge is "Cognitively accessing the fact that I am cognitively accessing the fact that I am currently exercising mastery of the sensorimotor contingencies of softness." Though defensible with the appropriate definition of "cognitive access", this usage has led to misunderstandings, because the higher and lower levels of cognitive access are different in the degree to which they are properly cognitive, and in the type of "latent" or "potential" facts that they concern. It would be interesting elsewhere to further clarify these issues, but here I wish to consider what seems like a much more important question.

5.3 Obtaining Phenomenality

Even if we accept (B) as a good approximation to *what it means* to be conscious of the softness of the sponge, there seems still to be a problem. Being conscious of the fact of softness is not the same as *consciously and phenomenally experiencing* the softness. Being conscious of a *fact* is phenomenally quite a different thing from *consciously feeling*. Feeling involves a phenomenal aspect, whereas the fact does not. How could having higher order cognitive access to a sensorimotor law provide an actual *feel*?

Indeed this is an issue that is skirted by HOT theories – these simply assert that "a mental state is conscious just in case it is accompanied by a higher-order thought to the effect that one is in that state" ([12] p. 352). From the sensorimotor point of view such a statement first is confusing, because mental states are not conscious, *people* are conscious. Second, while having such HOTs seems intuitively to be *what we mean* when we say we are conscious of a fact, something is missing when we are talking about a phenomenal experience, namely the experienced phenomenal quality of the presence of the environment.

The missing thing is provided by the sensorimotor approach. There are two ideas. The first depends on the fact that when you are cognitively accessing an ongoing sensorimotor interaction, you are not just accessing a single thing about that interaction. Among the battery of facts W, Y, Z... mentioned above and which are all *potentially* accessible, some contribute to providing an impression of "being connected to" or "in intimate contact" with the sensorimotor interaction that is going on. This is part of what provides the special impression of *phenomenal modality and presence*. The second idea concerns the special role that "grabbiness" plays in dominating our mental functioning, and thereby generating a form of *subservience* to our environments, which also strongly contributes to the impression of sensory *presence* or *reality*, of sensory experiences.

Phenomenal Modality and Presence. The main fact that you are cognitively accessing when you are experiencing the softness is that the current interaction obeys the laws of softness. However you are not simply cognitively accessing the fact of softness, which you could also get in a different fashion, say, by reading a book. There are additional facts available, such as for example that you are obtaining the feel of softness through a *tactile* interaction, and not through visual or auditory interactions - you have access to sensory modality ([13]). You can know this about your mode of access because you can access the fact that your finger motion is having an effect on sensory input. Another potentially available fact is that the interaction involves, say, your right hand, and these particular fingers. Such facts, of which there are very many, are not currently in the main focus of your attention . They are examples of the "latent" facts in the fringe of your awareness that you could, if you wanted, immediately make use of by bringing them into attentional focus. These potentially available facts are nevertheless part and parcel of the main fact, namely the fact of softness, and they contribute to the impression that the softness is of a type *obtained through touch*.

An even more important set of facts that you have potential access to are those that determine the impression of sensory *presence* of the softness. These facts are the bodiliness, insubordinateness and grabbiness of the ongoing interaction.

Bodiliness consists in the fact that moving your body (in particular your fingers) is what is currently providing changes in sensory input. This is a typical fact about sensory experiences *deriving from the outside world*, and it is not true of mental states like thoughts, memory or hallucinations, nor is it true of internal visceral states. Insubordinateness is also an indicator that you are *really currently interacting with the world*: it consists in the fact that although at this moment most of your changes in sensory input are coming from the changes caused by moving your fingers, there is an element of unpredictability: there is always the possibility that something might brush your fingers without you moving them (an insect inside the sponge?). Finally, and most important is that you have potential cognitive access to the grabbiness of the interaction. Grabbiness is the fact that it is possible that while you are engaged in exploring the sponge, your attention might be grabbed incontrovertibly by some sudden occurrence which is beyond your control: for example if you encounter a pin in the sponge that pricks your finger. Grabbiness is a fundamental property of sensory systems. Unlike thoughts and visceral states, sensory systems are hard-wired so that they can interrupt cognitive processing and orient attention to sudden sensory changes (see [1], [4]).

Bodiliness, insubordinateness and grabbiness are facts about your current sensorimotor interaction which are also part of the fringe of your higher order cognitive access to the softness. They are proofs that the current interaction is of an ongoing kind, that it is occurring with the outside world and is not purely a mental creation or some kind of visceral functioning. The potential to access these facts procures the feeling of modality (tactile, not visual or other), the “realness”, and the sensory “presence” of the interaction.

Presence and Subservience to the World. It is important to note the special role played by grabbiness. Contrary to what happens in normal thinking where our thoughts are more or less under our own control, in the case when we are involved in sensorimotor interactions with the world, because of grabbiness, our cognitive processes are partially *subservient* to the interactions. Because sensory systems are hard-wired to be able to interrupt our cognitive processing, we partially lose our mental control, and our attentional and cognitive capacities are drawn along and directed incontrovertibly by the ongoing interaction. Thus when we have higher order cognitive access to a sensorimotor interaction, not only are we aware of the battery of facts that correspond to the quality of that interaction, but our cognitive processes themselves are partially controlled and dominated by what is happening in that interaction. This potential direct effect of sensory stimulation on ongoing mental processing is an another important factor that gives sensory stimulation its specific feel of presence or reality.

In summary, when we have a sensory experience, the feeling we have of ongoing stimulation, of something occurring to us, of having an outside source of stimulation, can be subdivided into many sub-aspects such as those above. When I say it feels like something rather than nothing to have a sensory experience, this statement can be decomposed into being conscious of facts like: it’s modal (because it has the properties of touch, not of other senses), it has a particular quality (determined by the particular sensorimotor law), it’s real (because it has the bodiliness, insubordinateness and grabbiness of real interactions, and furthermore controls my thoughts the way real things do), etc. The claim is that there is nothing more to phenomenal consciousness than potential conscious access to all these facts.

Note defining conscious phenomenal experience as having higher order access to a sensorimotor interaction provides an improvement on HOT theories by giving an intuitively satisfactory account of the phenomenal aspect associated with conscious sensations. Of course in line with radical enactive approaches one could say that the "cognitive access" that I invoke does not have to be played out in a representationalist way. But some form of higher order "access" is necessary, simply because that is the normal usage of the word "conscious".

6 Conclusion

The sensorimotor approach is not a theory, but a way of being precise first about what we generally mean by the quality of an experience, and second about what we generally mean by being conscious of an experience.

As regards sensations, the approach proposes that experiencing a sensation involves being engaged in a sensorimotor interaction. The sensorimotor contingencies that describe that interaction constitute the experienced qualities of the experience. Far from "explaining away" the problem of qualia, the sensorimotor way of thinking provides immediate links between how we describe our sensory experiences (they have ineffable similarities and differences between their qualities, they are modal, they have sensory presence or realness...) and physical properties of the sensorimotor interactions involved (the sensorimotor laws, their bodiliness, insubordinateness and grabbiness...).

As regards what we mean by being conscious of something, the approach proposes that normal parlance requires appeal to a form of "higher order" cognitive access. This does not involve magical mechanisms and does not "sneak in" consciousness by the back door.

Unlike HOT approaches, the question of how higher order access to sensorimotor interactions can be accompanied by phenomenality receives an intuitive answer in terms of potential access to a battery of latent facts about the ongoing sensorimotor interactions, and in terms of the grabbiness of sensory systems, which creates a special subservience of mental states on the outside world.

Interestingly taking these definitions leads to some important theoretical, terminological and ethical questions. The ethical questions derive from the fact that the term "being conscious of something" demands the mental faculties needed for "higher order" access. If we want the term to retain the same meaning for different creatures, then we have to admit that consciousness, in this definition, must be present in progressively lesser degrees in children, in new-born babies and animals, whose mental faculties are less. This leads to the uncomfortable conclusion that creatures with lesser mental capacities do not consciously experience, for example, pain, in the same way human adults do, if at all. Trying to appeal to some other, "organismic" form of consciousness that all living creatures possess is not a defensible option, I have argued. The conclusion is that consciousness is a matter of degree, and so not an effective criterion to determine animal rights and whether we should give anesthetics to babies. We have to look elsewhere for reasons to justify our ethical practices.

Acknowledgements. I warmly thank Erik Myin, and particularly Jan Degenaar for help in clarifying this manuscript.

References

1. O'Regan, J.K.: Why red doesn't sound like a bell: Understanding the feel of consciousness. Oxford University Press, New York (2011)
2. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behav. Brain Sci. 24(5), 883–917 (2001)
3. Nagel, T.: What is it like to be a bat? Philos. Rev. 83, 435–456 (1974)
4. O'Regan, J.K., Myin, E., Noë, A.: Towards an analytic phenomenology: the concepts of 'bodiliness' and 'grabbiness'. In: Carsetti, A. (ed.) Proceedings of the International Colloquium: Seeing and Thinking. Reflections on Kanizsa's Studies in Visual Cognition, University Tor Vergata, Rome, June 8-9, pp. 103–114. Kluwer (2004)
5. Rosenthal, D.M.: From Rocco Gennaro (ed). In: Higher-Order Theories of Consciousness: An Anthology, pp. 17–44. Johns Benjamins Publishing Company, Amsterdam (2004)
6. Carruthers, P.: Higher-Order Theories of Consciousness. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy (Fall 2011 Edition) (2011)
7. Ryle, G.: The concept of mind. University of Chicago Press (1949)
8. Hutto, D.D., Myin, E.: Radicalizing enactivism: Basic minds without content. MIT Press, Cambridge (2012)
9. O'Regan, J.K., Block, N.: Discussion of J. Kevin O'Regan's 'Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness'. Rev. Philos. Psychol., 1–20 (2012)
10. Gallagher, S.: Philosophical conceptions of the self: implications for cognitive science. Trends Cogn. Sci. 4(1), 14–21 (2000)
11. Thompson, E.: Mind in life: Biology, phenomenology, and the sciences of mind, vol. xiv. Belknap Press/Harvard University Press, Cambridge (2007)
12. Rosenthal, D.M.: Consciousness and mind. Oxford University Press (2005)
13. Clark, A.: A Case where Access Implies Qualia. Analysis 60(265), 30–37 (2000)

Heideggerian Credentials? O'Regan's Sensorimotor Approach to Perception and Robots That Feel

Rachel Paine

Oxford University Department for Continuing Education, Oxford University, Oxford, UK
racheljpaine@yahoo.com

Abstract. Kevin O'Regan argues that seeing is a way of exploring the world, and that this approach helps us understand consciousness. O'Regan is interested in applying his ideas to the modeling of consciousness in robots. Hubert Dreyfus has raised a range of objections to traditional approaches to artificial intelligence, based on his reading of Heidegger. In light of this, I explore here ways in which O'Regan's approach meets these Heideggerian considerations, and ways in which his account is more Heideggerian than that of Dreyfus. Despite these successes, O'Regan leaves out any role for emotion. This is an area where a Heideggerian perspective may offer useful insights into what more is needed for the sense of self O'Regan includes in his account in order for a robot to feel.

1 Introduction

Kevin O'Regan argues, in support of his sensorimotor approach to perception (O'Regan 2011), that vision is a way of manipulating the environment, an exploratory activity, one motivated and sustained by our interest in our world. Perceptual experience is not *generated* by brains, but *is constituted by* our uses of our perceptual systems. A central conclusion he draws from the empirical data he has accumulated over several decades is that there is nothing intrinsic to the way our sensory systems are set up that explains the experience we have of a continuous and coherent world. This continuity and coherence is present to us because our activities presuppose continuity and coherence.

This understanding of perception as engagement rather than representation has been around for a while, as O'Regan notes: in 1962 he heard Donald M. MacKay, a distinguished neuroscientist publishing on perception in the 60s, 70s and 80s, give a lecture about perception, in which he claimed that "the eye was like a giant hand that samples the outside world." (O'Regan 2011: 23) and Merleau-Ponty held that vision "was a form of palpation." (Ibid: 23) Varela, Thompson and Rosch, in *The Embodied Mind* (1991) argue that objects are not seen by our extracting their features in order to construct representations of them, but by our exploratory activities. Empirical evidence suggests that not only is active exploration of the world needed for the development of perceptual experience, but such experience also depends upon the

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*, Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_3, © Springer International Publishing Switzerland 2014

regularity of this activity.[1] This historical layering of experience can be thought of as a form of sedimentation, establishing what O'Regan refers to as having a grasp of the "sensorimotor contingencies", or "laws", the ways in which perceptual experience is determined by experiences formed over time.

Much of the philosophical work that supports the view that our engagements constitute our perceptual experience comes from philosophical phenomenology. Merleau-Ponty developed a detailed, embodied phenomenology of perception, bringing in the centrality of the body (what O'Regan refers to as the bodily condition) for sensory experience, and introduced the notion of the "intentional arc" (Merleau-Ponty 2003) referenced in Hubert Dreyfus's arguments for an embodied AI. In terms of critiques of AI in its more traditional forms, both Dreyfus's and John Haugeland's work have stood on the shoulders of Heidegger.[2] Since O'Regan has chosen as his final point, indeed the very last word in his 2011 book, to consider the implications of his approch for robotic consciousness, I will conclude with a discussion of some relevant Heideggerian concerns for this account as a model for consciousness.

The paper has three parts: in part 1 I explore the ways in which O'Regan's account satisfies Dreyfus's Heideggerian considerations. In part 2 I argue that O'Regan is, in fact, more Heideggerian than is Dreyfus. O'Regan goes further than Dreyfus in his recognition of the importance of our interests and contexts in constituting sensory experience. His discussion of the role of the self also brings him closer to Heidegger than Dreyfus. In part 3 I raise a concern about O'Regan's account of the emotions. What is missing in his account is the centrality of the emotions required for the self to play the grounding role O'Regan attributes to it.

2 Phenomenological Support for O'Regan's Sensorimotor Approach

2.1 The Intimacy of the Mind, Body and World

John Haugeland has been a key philosopher challenging "good old fashioned AI (GOFAI)" (Haugeland 1985), the traditional view of mind as disembodied computational systems. Haugeland argues that there is no inner quality of consciousness that shows up *in addition to* the qualities of being "intimately" in a world. Haugeland defines the concept of intimacy this way:

> The term *intimacy* is meant to suggest more than just a necessary interrelation or interdependence but a kind of *commingling* or *integralness* of mind, body, and the world – that is, to undermine their very distinctness. (Haugeland 1998:208)

[1] Varela, Thompson, Rosch (1991) for a discussion of Walter Freeman's experiments involving the olfactory experiences of rabbits. Freeman concluded that rabbits did not perceive sensory stimulations until these had become regular occurrences, suggesting a need for these to appear as parts of practices or habits of experience in order for them to be perceived at all.

[2] See Froese, T., and Ziemke, T. for cognitivist scientists' views on the importance of Dreyfus's Heideggerian critique.

The concept of intimacy is meant to capture the dynamic interaction of a creature in its world, its "embodiment and embeddedness in the world." (Ibid: 208). O'Regan argues from his understanding of our perceptual systems to the effect that these are not adequate for the perception we experience without our being in a world in this *commingled* fashion. To make his case, O'Regan outlines the inadequacies of the visual system in delivering anything like a representation of the visual scene. For instance, there is a "blind spot" where the optic nerve emerges from the eye, taking up the space where photoreceptors (necessary for registering light) would otherwise be. The photoreceptors themselves are distributed unevenly, concentrated toward the centre of the eye, so that what can be seen peripherally is significantly less detailed. This lack of detail also affects the depths of the colours we see. The phenomenon of "cortical magnification", where the retinal neurons bring together the reflections of light off an object, produces further distortions. O'Regan argues that none of the standard explanations for how we 'fix' the images received in order to experience them as veridical is successful. The view that perception is dependent on "compensatory mechanisms" to improve the internal representations implicitly assumes that there is a "homunculus" that can adjust what is received to what is real, comparing the internal image with the external reality. Such mechanisms for filling in blind spots and adjusting blurred presentations, if they did exist, turn out, in fact, not to do a very good job of it. We, in fact, "see" in a strictly technical manner pretty inaccurately what is presented to us.[3] Instead, O'Regan argues that we experience the world perceptually in virtue of our interactions with it. We are not attempting to produce a veridical representation, but are engaged in activities that produce our visual experience of the world. We see the world as containing objects whose shapes we understand through past experience of them and current interest in them. We experience visual scenes and objects as real, three dimensional and familiar because we know we can explore them and have seen them from other perspectives before. O'Regan refers to our experiences as being "at home" with the objects or places that make up our visual landscape. Prior experiences of these have built up a grasp of the world we are confident can be explored in particular ways. We do not need to refer to an inner representation, since, as Rodney Brooks argued, "the world is its own best model." (Brooks 1990)

O'Regan's conclusions are that the close coupling of the organism's sensory apparatus with actions within the world constitute qualitative experience through "skilled modes of interaction with the environment." (O'Regan 2011: 115) The "skills" involved here are laid down through the history of interactions sedimented through our repeated interests and needs and the contexts in which these are expressed. Similar to the notion of skill at playing tennis, these skills are a habituated, embodied grasp of how to engage in dynamic situations.

Breaking down this interactive account of perception, O'Regan finds that it can be reduced to four qualities of experience:

---

[3] See chapter one of O'Regan (2011) for the full description of the phenomena, the theories and their shortcomings.

Richness: Our grasp of the world has an open-ended quality to it… there is always more to it to be explored. Perhaps this quality of richness is what gives us the experience of being "at home" in a world that extends infinitely outward.[4]

Bodiliness: Our bodily interactions are a constituent of our experience of the world. In this way we grasp as fundamental that we are in-a-world, not observers of it, as we might experience ourselves when watching a distant scene played out on a screen.

(Partial) Insubordinateness: The world is not shaped by us, but partially imposes itself on us, so giving us a grasp of its extending beyond us.

Grabbiness: We are set up to be drawn to sudden changes in our environment, from loud noises to flitting movements. "Grabbiness" suggests at least a minimal form of desire is operative: we are not impartial to what goes on around us, but are motivated to want to know what's out there, what has changed, what might be dangerous or useful. "Grabbiness" seems tightly connected to a nature motivated to be concerned about its world.

Mere memory does not have these qualities to any significant degree. The presence of these gives our experience its "feel" of reality, or "presence".

O'Regan's (2011) position is developed through both positive and negative support: on the positive side, he finds that our "probing" of the environment produces our experience of it. If we are interested in it, then we seek it out, and are attuned to finding it. If we do not, then we tend to miss it. On the negative side, he argues that what our sensory systems actually deliver is so deficient in content that if we did rely on representations caused by the sensory systems themselves, we would have a very patchy and unstable view of the world. Instead, given the poverty of the stimulus, we, nonetheless, experience the world as coherent and consistent. We do this because this experience is constituted by our interactions, not by images generated in the brain.

2.2 Dreyfus's Heideggerian Conditions for Human Mentality

Hubert Dreyfus argues that considering a range of Heideggerian insights into how we engage with the world reveals what classical AI, or the computational model of mind, misses. One such insight is found in Heidegger's account of the objects we encounter in our daily lives as "equipment" or that which is "ready to hand". The idea behind that of "readiness-to-hand" is that objects are relevant to our projects and intentions, and have their meanings determined by the contexts of use.

Dreyfus refers to Heidegger's account of "ready-to-hand" objects as constituents of our needs and interests:

> Heidegger describes our most basic experience of what he later calls "pressing into possibilities" not as dealing with the desk, the door, the lamp, the chair and so forth, but as directly responding to a "what for":

[4] O'Regan has recently dismissed the quality of richness as central to experience since it appears in non-conscious experience also. I would like to include it, however, as perhaps straddling both the conscious and unconscious, as, in fact, do the other qualities he discusses.

> 'What is first of all 'given'… is the 'for writing', the 'for going in and out,' the 'for illuminating,' the 'for sitting.' That is, writing, going-in and-out, sitting and the like are what we are a priori involved with. What we know when we 'know our way around and what we learn are these 'for-whats.' (Dreyfus 2007, 252)

Equipment is what it is only in virtue of the uses to which we put it. I need to know what a bicycle is for if I am to ride it. I need to know what a tennis racquet is for if I am to swing it appropriately. My experience of engagements with the objects of my practical daily life is constituted by my background grasp of their 'for-what". This resonates with Gibsonian ideas of affordances: no object is a mere object, but is experienced as for-sitting-on, or drinking-from. Our perceptual engagement of these objects necessarily makes the opportunities or possibilities that they afford us the focus for what qualities of these objects we experience. The idea that the "for-what" of an object is the way in which we experience it should not suggest that "for-whats" inhere in the objects. A hammer used for drawing a line in the sand becomes ready-to-hand as a line-drawer, rather than a pounder of nails. "for-whats" arise from the interactions between the objects and the user, another way in which "intimacy" is a condition for our experiences.

For artificial intelligence to approach human mentality, our robots must have the ability to immerse themselves in the world of daily practices with these objects, such that the objects of the robot's world are ready-to-hand, reflecting the interests and concerns of the robot in its use of them. An artificial intelligence system must be relevantly situated in a world in which the objects with which it engages are meaningful to it. The problem facing a programmed computational system is that such programming will not provide the system with the flexibility required to have a genuine engagement with its environment. The alternative is a system that is coupled with the environment. Rodney Brooks explains:

> Nouvelle AI is based on the physical grounding hypothesis. This hypothesis states that to build a system that is intelligent it is necessary to have its representations grounded in the physical world. Our experience with this approach is that once this commitment is made, the need for traditional symbolic representations soon fades entirely. The key observation is that the world is its own best model. It is always exactly up to date. It always contains every detail there is to be known. The trick is to sense it appropriately and often enough.

To build a system based on the physical grounding hypothesis it is necessary to connect it to the world via a set of sensors and actuators. Typed input and output are no longer of interest. They are not physically grounded.

…

This suggests that problem solving behavior, language, expert knowledge and application, and reason, are all rather simple once the essence of being and reacting are available. That essence is the ability to move around in a dynamic environment, sensing the surroundings to a degree sufficient to achieve the necessary maintenance of life and reproduction. (1990: 6)

The system will need to be designed so that it is receptive to interactions with the environment rather than programmed to control it. This is an idea that is captured well in O'Regan's experimental work. We do not perceive objects or a visual scene in an objective way, as a view from nowhere or as a view without intention. Our history of engagements with the world's objects as we have an interest and need for them lays down the sensorimotor contingencies that constitute experience. In a passage that resonates with Heideggerian considerations, O'Regan describes the role of this historically-laden understanding of ourselves in our world that is the basis for grasping the world, rather than any necessity for acting at the moment:

> The idea is similar to the idea of feeling at home. When I am sitting on my sofa, I feel at home because there are a variety of actions I can undertake (go into the kitchen and get a coffee, to the bedroom and lie down, etc.). But I need not undertake them. It is because I am poised to do these things that I have the feeling of being at home. Feeling at home does not require actual action. In the same way seeing does not require actual action. It requires having previously acted, and it requires having the future potential for action. (O'Regan 2011: 87)

"Being at home" can be understood in terms of Heidegger's notion of the "ready-to-hand", since it is being in a world we grasp in virtue of its temporal extension, its embeddedness in cultural values, and its containing within it the possibilities for current and future actions. It is part of a web of interactions, in a world that shapes our understanding of its artifacts, its customs, and, in this example, the need and desire for shelter, the "for-whats" that constitute experience.

John Haugeland also discusses the centrality to experience of engagements that are the result of past experiences sedimenting into a present understood in terms of possibilities in his Heideggerian insights into the role of culture and practice:

> Human intelligence is surely manifested in the ability to design and make things—using, as the case may be, boards and nails. Now, for such a design to work, it must be possible to drive nails into pieces of wood in a way that will hold them together. But neither a designer nor a carpenter ever needs to think about that—it need never even occur to them. (They take it for granted, as a fish does water.) The suitability of these materials and techniques is embedded in the structure of their culture: the logging industry, the manufacture of wire, the existence of lumber yards—and, of course, countless bodily skills and habits passed down from generation to generation. (Haugeland, 1997: 26)

Haugeland's description of the engaged practices here reflects what O'Regan is also concerned with: there is a background of experience that has a structure arising from the temporally-extended grasp of our environment and our practices within it that makes what we are doing and perceiving meaningful. This meaningfulness is not something over and above the practices themselves, but is the quality of experiencing a direct engagement with the world.

2.3 The Frame Problem

For my purposes here, I will view the frame problem as the problem of how a system might choose what is relevant to it, given an infinite amount of data. For the simple exercise of playing chess, a system needs to know that winning is the objective, and playing by the rules is the only means available, as well as knowing the rules. Clearly, in the case of human mental life, our engagements are much more open ended than a game of chess. Even a game of tennis involves recognition of complex interactions with others that are difficult to specify. So the frame problem remains a stumbling block for developing a system that requires context sensitivity for it to work out how to proceed. The idea that objects play a role in a temporally extended world of ongoing engagements is only the first step. The system needs to be so immersed in that world that the role these objects play can determine their value to us. A hammer is only meaningful to us within the context of its uses and our needs for its uses. That we have needs and desires is central to this Heideggerian approach to understanding objects and their (and our) world. If this understanding is in place, then there should be no frame problem to solve. To arrive at this dissolution of the frame problem, Dreyfus turns to Merleau-Ponty. Dreyfus describes Merleau-Ponty's intentional arc as the gestalt or unity of the world insofar as it is "organized in terms of an organism's need to find its way around." (Dreyfus 2007: 255) He says:

> ...in our skilled activity we are drawn to move so as to achieve a better and better grip on our situation..... acting is experienced as a steady flow of skillful activity in response to the situation. One does not need to know what the optimum is in order to move towards it. One's body is simply drawn to lower the tension." (Ibid, 255)

For Dreyfus, the intentional arc is the set of conditions the world offers us for satisfying our needs and desires. The needs and desires involved are the basic kind facing any creature that has the autonomy to fend for itself, while the environment is that which puts pressure on the creature or offers it opportunities. The creature, meanwhile, responds in whatever way will maximize its grip on the world and find equilibrium. Dreyfus takes Merleau-Ponty to be describing a "feedback loop" between the organism and the perceptual world. While this may be true, the intentional arc Merleau-Ponty describes references a much more complex world, one that reaches beyond the satisfaction of basic needs.

Merleau-Ponty describes the intentional arc:

> “The life of consciousness - Cognitive life, the life of desire or perceptual life - …is subtended by an "intentional arc" which projects round about us our past, our future, our human setting, our physical, ideological and moral situation, or rather results in our being situated in all these respects. It is this intentional arc which brings about the unity of the senses, of intelligence, of sensibility and motility." (Merleau-Ponty 2003: 157)

The "intentional arc" refers to the temporal, spatial, and cultural totality within which we grasp our world as meaningful. Merleau-Ponty emphasizes the cultural by breaking this down into our "ideological" and our "moral" situation. The political and moral frameworks that establish values are integral to the intentional arc. Merleau-Ponty's conclusion is that we cannot give an account of the senses, of our actions, or of our thoughts, without reference to their unity within this greater, temporally-extended and socially-constituted world. We are not to be viewed as passive recipients of this world, but as constituents of it. This would suggest that it is our experience embracing these ideas and values that establish us as having a world. Having a world is just experiencing the unity of the being who senses, thinks, and acts in that world. That we are interested in our world is fundamental to this conception of our being immersed within an "intentional arc", echoing Heidegger on our *care* for the world as the ground of our being-in-the-world.[5]

The intentional arc takes us far beyond the feedback loop of an organism coupled with its environment for the purposes of satisfying its needs. The intentional arc makes the question of the frame problem irrelevant. Despite the limited use to which Dreyfus puts this rich concept, he is right that a description of how we do engage with the world that includes an intentional arc does dissolve the frame problem. We are in a world already, we do not have to define it. The world is presupposed. The frame problem does not arise with O'Regan's approach either, for reasons that resonate with the ideas of Merleau-Ponty here. We are already 'in-a-world' such that how we see things is already framed.

3 O'Regan Is More Heideggerian than Dreyfus

Central to O'Regan's account is that there is an "I" who experiences the three or four qualities of experience summarized above. In this he is closer to Heidegger's own understanding than Dreyfus. Dreyfus says that the "I" disappears into the activities:

> When immersed in the world of daily coping, "normally there is no "I" and no experiencing of the door at all but simply pressing into the possibility of going out.... there is no experience of an entity doing the soliciting; just the immediate response to a solicitation. (Dreyfus 2007: 252)

Dasein, Heidegger's term for human being, does not ever stop being an entity for itself. In fact, it is the concern Dasein has primordially with itself (always in a world) that leads it to use an object with a purpose. Heidegger makes this point in reference to hammering:

---

[5] Heidegger's account of the surrounding world, *umwelt*, the world with others, *mitwelt*, and self-world, *selbstwelt* and their convergences is explained well in Scott Campbell (2012). Altogether these constitute *sorgenwelten*, or the care-world.

> With the "towards which" of serviceability there can again be an involvement: with this thing, for instance, which is ready-to-hand and which we accordingly call a "hammer", there is an involvement in hammering; with hammering there is an involvement in making something fast: with making something fast, there is an involvement in protection against bad weather; and this protection "is" for the sake of providing shelter for Dasein - that is to say, for the sake of a possibility of Dasein's Being. ... the primary "towards- which" is a "for-the-sake-of-which". But the "for-the-sake-of" always pertains to the Being of Dasein, for which, in its Being, that very Being is essentially an *issue*. (Heidegger 1962: 116 - 117)

Dasein is not absent in its concernful practices, but is disclosed by them. Dasein's openness to the world is not a dissolution of itself, but a constituent of itself. Only in the sense of Dasein losing its way in the everyday, does Dasein lose itself. But Dasein's engagement with that which is ready-to-hand does not represent the fleeing from the world that Heidegger describes in that case, but, rather, the openness of Dasein to its world by caring for both itself and its world. Nor is that world just made up of the equipment that we use in order to work at something for Dasein's sake: that equipment itself has meaning in the context of Dasein's embeddedness in its world:

> But the work to be produced is not merely usable for something. The production itself is a using of something for something. In the work there is also a reference or assignment to 'materials': the work is dependent on leather, thread, needles, and the like. Leather, moreover is produced from hides. These are taken from animals, which someone else has raised. …hammer, tongs, and needle, refer in themselves to steel, iron, metal, mineral, wood, in that they consist of these...... (Heidegger 1962: 70-71)

Every practice is part of a web of interactions that constitute the world of Dasein. This is the world that discloses Dasein, that lets Dasein be Dasein itself.

O'Regan's approach, in which our perceptual engagements over time lay down or sediment the structure of experience, acknowledges that these past engagements are culture bound:

> Social psychologists studying the unconscious influence of cultural prototypes on our behavior show that our everyday actions are more determined than we think by automatic, socially-driven influences. We unconsciously espouse images of ourselves as having a certain personality, as belonging to a particular social category, and these cultural prototypes strongly influence the construction of our identity. Indeed, a person's gait, gestures, speech, taste, and dress are all exquisitely sensitive to their cultural or social context. (O'Regan 2011: 82)

We are shaped by and continue to shape ourselves by reference to a greater culture: what interests us and what contexts we have experienced come from the particular world in which we find ourselves over the course of our lifetime.

3.1 Sensorimotor Theory and Consciousness

O'Regan's account takes consciousness to be constituted by engagements that have particular qualities, those of richness, bodiliness, partial insubordination and grabbiness. Focusing on these qualities allows us to make sense of the differences in conscious experience between slugs, human babies, and adult human beings. The minimal cognitive access of a human infant is qualitatively different from that of an adult. O'Regan takes this difference to involve higher-order awareness:

> Having conscious access involves not only cognitively accessing something in order to exercise a choice about what to do with respect to that thing but also being aware of the whole context within which you are doing that cognitive accessing. Thus, it involves being ready to show, by choosing from a wider set of alternative actions, that you can make use of your awareness that the more restricted cognitive accessing is going on. (O'Regan 2011: 91)

With cognitive access referring only to the function of making choices among options, something chess playing computers can do, "consciousness" is understood as referring to grasping the context within which these choices are made, thus extending our range of choices. If higher-order access is in place, a chess playing machine would then make choices that involve consciousness:

> This carries the implicit assumption that there are a variety of possible other things that the machine could have been poised to make use of, like the expression on your face, for example, or the fact that it's playing chess and not dominoes. The second context, possibilities for action, derives from the variety of things the machine could do about the fact that it is poised to make use of your moves (It could carry on playing, but it can also do other things, like talk about your move or ignore your move and talk about the weather). (Ibid: 91)

If a machine can make use of expressions, and choose to talk about the weather, this would suggest that the range of what the machine might choose to do is not something programmed within it. What might be the source of these choices? O'Regan suggests that it might be the "self", the centre of concern missing from Dreyfus's account, and central to a Heideggerian conception of our mental life. With the higher order awareness in place, the machine now has a first-person perspective. O'Regan explains:

> The machine is not only poised to apply its cognitive abilities, but it also knows that *it* is in this way poised. … Furthermore, if its self is well developed, it might also know that this entity can be considered by others, and by itself, as having desires, motivations, purposes, plans, reasons, and intentions. (Ibid: 92)

Not only does the machine self have its own internal milieu, in which desires, motivations, and the rest are there, guiding action, but it sees itself as a socially-embedded being:

> Such a self is socially defined, in the sense that it knows the social presuppositions and implications of the situation (the fact that it is supposedly good for it to win the game, that presumably you yourself also want to win, that you and it have come together and agree to obey the rules…) (Ibid: 92)

O'Regan's proposal is that this awareness is a necessary condition for consciousness. The minimal sense of self other animals and babies have means that, despite the presence of *raw feels,* their consciousness is not like ours. In the case of other animals and young infants, O'Regan says:

> ...its organism is undoubtedly reacting in response to sensory stimulation, but there is not very much of a self for the organism to feel it, at least not in the way adult humans feel….. in the case of pain, the organism is providing an avoidance reaction, registering a stress response, signaling by its crying that it requires help from its conspecifics. But since there is no structured "I" to know and cognitively use the fact that these things are going on in the body, we logically cannot say, …. that the animal or baby, considered as a "self" feels anything in the same way as adult humans feel it. (Ibid, 123)

O'Regan's description of the higher-order awareness involved in adult consciousness implies that, along with the higher-order awareness, there is a sense of agency. A creature could be aware of its states, and of its perspective, without having agency. It is the presence of agency, of being not only self-aware but also *self-motivated*, *choosing* options and *embracing or modifying* received norms that is a pre-requisite for consciousness as we are describing it here.

3.2 Agency

When you are paying attention to something, you can miss what else is going on. For example, in the midst of scoring against the other team in a game of football, a player won't notice the particular people cheering him on. Once his team has won the game, and faces the cheering fans, however, the crowd will fill the stadium in sudden technicolour. As shown in the empirical work done by O'Regan and many others, we can miss experiences we are not focusing on if they have little or no significance for us.

Having a focus requires engaging with the world through a perspective structured by our concerns and interests, within contexts. O'Regan argues that awareness must be awareness "of" *and* awareness "through" the self. A camera is aware "of" the objects in front of it, it is not aware "through" its own interests. Any creature without this self-perspectival awareness therefore, has awareness "of", in the minimal sense of being directed toward the environment, CCTV style, but not awareness "through" the filter of its own interests. This distinction helps clarify the role higher-order thought plays in O'Regan's account: the point at which we can talk of consciousness being constituted by sensory engagements is the point where these engagements are structured, not only by the temporal layering or sedimenting of experience but by the creature's possession of a self-directed concern for itself-in-its-world.

The roles objects have in our engagements are determined both culturally and through our own self-directed concerns and interests. What we ask of them and expect of them determines how they show up for us.

John Dewey made a similar point in 1929 when he described how the character of an object is necessarily determined by one's grasp of its role in one's life:

> Meanings acquired in connection with the use of tools and of language exercise a profound influence upon organic feelings. In the reckoning of this account, are included the changes effected by all the consequences of attitude and habit due to all the consequences of tools and language – in short, civilization… The subconscious of a civilized adult reflects all the habits he has acquired; that is to say, all the organic modifications he has undergone. (Dewey 1929: 300)

This description of what Merleau-Ponty would later refer to as the intentional arc, that which subtends all perception, is given fuller detail in Dewey's illustration of the role of function and understanding in our perception of objects:

> The same existential events are capable of an infinite number of meanings. Thus an existence identified as "paper", because the meaning uppermost at the moment is "something to be written upon," has as many other explicit meanings as it has important consequences recognized in the various connective interactions into which it enters. Since possibilities of conjunction are endless, and since the consequences of any of them may at some time be significant, its potential meanings are endless. It signifies something to start a fire with; something like snow; made of wood-pulp; manufactured for profit; property in the legal sense; a definite combination illustrative of certain principles of chemical science; an article the invention of which has made a tremendous difference in human history, and so on indefinitely. (Ibid, 319-320)

The crucial point is that the associations are indefinite, infinite. You cannot program something to have all of these in mind, because you cannot program an infinite number of possibilities. More interestingly, you cannot program knowledge of these possibilities because we have an implicit, not explicit, grasp of the possibilities. Only if we have some motivation already established by the history of our concerns and interests can we then perceive something in a particular way. This is a point that O'Regan makes in his account of perceptual selection. We draw on previously laid down experiences and, in combination with our current interests, we shape what is before us into a perception that "makes sense'. This explains why, although what is given to us merely physiologically is incomplete, we are able to experience a coherently-perceived world. We are capable of so structuring our experience because we have seen these things before in this way and we are looking for them to be present to us in this way now.

In summary, O'Regan's account resonates with a range of phenomenologically-grounded accounts. His inclusion of the self as the structuring viewpoint through which the world is perceived as making sense has not played a sufficiently central role in AI, suggesting that these accounts miss something necessary in the conditions required for consciousness.

4 Where O'Regan Might Develop Heideggerian Insights Further

4.1 Emotion

One problem with O'Regan's account is in the role given to emotion[6]. On this approach, our emotions are just one experience in the array of possible experiences and are not essential to the self that O'Regan argues constitutes, along with higher-order awareness, adult human consciousness.

O'Regan first mentions emotions in a passage describing what aspects of our experiences are non-essential, or "*add-ons*" to the consciousness experienced. "Emotions like fear, anger, and shame...would appear to involve specific bodily manifestations such as changes in heartbeat, flushing, or other reactions of the autonomic nervous system." (O'Regan 2011: 95) He concurs with a scientific view that the feeling of the emotion lies in the higher order awareness of the bodily changes, a view that is in part that of William James, and has currency today. On this understanding of the emotions, the sensorimotor view might be well placed to explore the idea that we have a higher-order awareness of our bodily changes that extends to include a higher order awareness of the ongoing interactions with the environment that constitute conscious experience.

But here O'Regan distinguishes the higher-order awareness of emotions from that of sensory experience. Although similar, there are crucial differences:

> Certainly emotions have grabbiness. Fear, if it is the prototypical emotion, may completely grab your attentional and cognitive resources and prevent you from functioning normally. The grabbiness is perhaps slightly different from the grabbiness of a sensory feel, because the grabbiness of fear requires a cognitive and probably a social interpretation. (O'Regan 2011:170)

[6] There are many accounts of the emotions that draw distinctions between the concepts of "emotion", "affect", and "mood". My own understanding of these terms is that these distinctions may be instrumentally useful, but are not distinct categories in themselves. I use the term emotion when O'Regan, or James does, or the term "mood" when referring to Heidegger. However, "affect" adequately covers the range of feelings. Even cognitivists such as Richard Lazarus (1991) acknowledge that our emotions include the activity of the viscera to which William James points when describing the emotions experienced as bodily feelings. (1889)

Because fear requires cognition it is not as strong as the grabbiness of sensory perception[7]. Similarly for the bodiliness of emotions in general: they are not as phenomenally present as sensory perceptions. For this reason, they are really just our higher-order awareness of bodily states (plus the cognitive and socially-mediated interpretations we give to them).

So, unlike sensory perception, in which our temporally-extended, socially-interpreted habits of interaction give rise to a high level of phenomenal presence, our emotions are just about bodily states we are currently experiencing, with the historical and socially-interpreted habits of interaction *added on* to the experience of passing sensations in the body.

In order for a being to have an immersed interest in her world, she must have, at the ground of this interest, an emotional, affective nature. This is a central insight of Heidegger's, so I will say something about what this affective account offers us here. Heidegger takes mood to be the ground of Being: "…ontologically mood is a primordial kind of Being for Dasein, in which Dasein is disclosed to itself *prior* to all cognition and volition…" (Heidegger 1962: 175) It is through our moods that we are disclosed as ourselves. Mood's nature as ontologically prior to all else cannot be overemphasized: when experiencing ourselves as emotional beings, we are experiencing ourselves as ourselves, not as a neutral self having an additional experience of some emotional kind:

> [A mood] comes neither from 'outside' nor from 'inside' but arises out of Being-in-the-world, as a way of such Being…. The mood has already disclosed, in every case, Being-in-the-world as a whole, and makes it possible first of all to direct oneself toward something. (Ibid, 176)

Mood is fundamental to our being, prior to everything else and necessary for any interest in the world to take place. Without mood, there can be no directedness upon the world, since Dasein's acting in the world always reflects its care towards itself in that world. Mood grounds our experience in the world as that of our own, that which the *self* dynamically produces.

Without this priority placed on mood, care itself could not have the grounding role Heidegger gives it, and which is implicit in any account that takes "interest" or "focus" to be motivational.

Heidegger's account of a contextualized world in which we engage for the sake of our *concerns* and through our interests (the "seeing through" that makes for the agential self) provides philosophical support for approaching perceptual experience as sensorimotor laws laid down though past experience. The separation of the emotional life from the rest of the temporally-extended and contextualized experience actually severs the connection between "mere" perceiving and the perception that involves our interests and concerns.

O'Regan has gone part way in establishing a role for the "self", for whom the perceptual experiences have meaning gained through the engagements with the world of a socially-embedded human being. However, grounding our *interest* in

[7] The work of LeDoux (1996) and others confirms the presence of fear in the absence of cognitive input.

experiencing the world, is the *care* for ourselves in that world. This care needs to be seen as the affective, motivational ground of the self.

There is an abundance of support for such a view. For instance, an account of the motivational nature of our emotions as fundamental to experience has been developed by Jaak Panksepp, a neurobiologist focusing on the neurobiology of our emotions. Panksepp joins the ranks of a number of researchers who view the emotion systems as fundamental to any account of a self-motivated perspective. These affects are seen to be the source of the very experience of being an "I" (Panksepp 1998, Damasio 2000, Stern 2000) When O'Regan cites "interest" in something as a basis for our focusing on and experiencing it, or lack of interest as a basis for our not experiencing it, he is assuming the presence of an affective nature structured to care about the world, to have interests, concerns and aims, as the basis or constituents of the self.

5 Conclusion

O'Regan's account of the sensorimotor contingencies that sediment to constitute the qualities of experience is one with philosophical predecessors and rich support from empirical science. More than that, it fares far better as a resource for contemporary work in AI, than does traditional AI in reflecting the work done in embodied and embedded consciousness. Where I think O'Regan misses a central feature of the conditions of consciousness is in what he has to say about our emotional life. The emotions are not just a feature that we experience bodily, with additional cognitive interpretation. Our emotional nature is the ground of the interest-focused self that makes our consciousness that of a being historically and socially embedded, with a self-awareness that is experienced as an agential perspective on the world.

References

1. Brooks, R.: Elephants Don't Play Chess in Robotics and Autonomous Systems 6 (190), 3–15 (1990)
2. Campbell, S.: The Early Heidegger's Philosophy of Life: Facticity, Being and Language. Fordham University Press, New York (2012)
3. Damasio, A.: The Feeling of What Happens. William Heinemann, London (2000)
4. Dewey, J.: Experience and Nature. W.W. Norton & Co., New York (1929)
5. Dreyfus, H.L.: Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian. Philosophical Psychology 20(2), 247–268 (2007)
6. Froese, T., Ziemke, T.: Enactive Artificial Intelligence: Investigating the Systemic Organization of Life and Mind. Artificial Intelligence 173(3-4), 466–500 (2009)
7. Haugeland, J.: Artificial Intelligence: The Very Idea. MIT Press, MA (1985)
8. Haugeland, J. (ed.): Mind Design II. MIT Press, MA (1997)
9. Haugeland, J.: Having Thought. Harvard University Press, MA (1998)
10. Heidegger, M.: Being and Time. Macquarrie, J., Robinson, E. (trans.). Harper-Collins, San Francisco (1962)
11. James, W.: What is an Emotion? Mind 9(34), 188–205 (1884)
12. Lazarus, R.S.: Emotion and Adaptation

13. LeDoux, J.: The Emotional Brain: The Mysterious Underpinnings of Emotional Life. Simon & Schuster, New York (1996)
14. McLuhan, M.: Gutenberg Galaxy. Routledge and Kegan Paul PLC, London (1967)
15. Merleau-Ponty, M.: Phenomenology of Perception. Smith, C. (trans.). Routledge, London (2002)
16. O'Regan, J.: Why Red Doesn't Sound Like a Bell. OUP, New York (2011)
17. Panksepp, J.: Affective Neuroscience. OUP, Oxford (1998)
18. Stern, D.: The Interpersonal World of the Infant. Basic Books, New York (2000)
19. Varela, F., Thompson, E., Rosch, E.: The Embodied Mind. MIT Press, Cambridge (1993)

The Phenomenology of Sensorimotor Understanding

Ken Pepper

Department of Philosophy, University of York, UK

Abstract. This paper draws on Maurice Merleau-Ponty's philosophy to sketch a phenomenological interpretation of sensorimotor understanding. I begin by situating Noë's enactive theory of vision in relation to Husserlian phenomenology. I then raise three related objections to Noë's treatment of sensorimotor understanding in terms of practical knowledge of possibilities for action. Finally, I appeal to *Phenomenology of Perception* to show how two of its major operative concepts – the 'body schema' and 'sedimentation' – can help to plug the gaps in Noë's account.

Keywords: Merleau-Ponty, Alva Noë, body schema, affordances, sensorimotor enactivism.

1 Introduction

Although Merleau-Ponty's early phenomenology is often invoked in discussions of embodied and enactive approaches 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 enactive approach, without due attention to the wider philosophical project in which those descriptions occur. 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 contemporary sensorimotor theorists and phenomenologists. This paper sketches one path a more satisfying engagement might take, by casting a Merleau-Pontian eye over the role of so-called 'sensorimotor understanding' in visual experience. I begin by introducing O'Regan and Noë's notion of sensorimotor understanding and situating it relation to Merleau-Ponty's own philosophical starting point, Husserlian phenomenology. I then present three problems incurred by Noë's characterisation of sensorimotor understanding in terms of practical knowledge of possibilities for action. The decision to focus primarily on the work of Noë rather than O'Regan is motivated by Noë's repeated assertion that his philosophical project is essentially a phenomenological one [e.g. 2004 p.33; 176]. I then turn to *Phenomenology of Perception* to show how two of its major themes – the 'body schema' and the arguably lesser known 'sedimentation' – can help plug the gaps in Noë's account and form the basis of what Husserl called a 'genetic' phenomenology of sensorimotor understanding.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_4,

2 Object Horizons, Affordances, and Sensorimotor Understanding

It is tempting to think that only that which imposes on the retinas can be presented in visual experience, but this commits what Merleau-Ponty calls the 'experience error' – the (mis)description of perception in terms of what we come to know upon reflection about its objects at the expense of capturing the original lived experience [2012 p.5]. In fact, though strictly speaking unseen, occluded surfaces and features of objects are 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. 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. Husserl revisited these phenomena of perspective and presence throughout the development of his phenomenology. In Husserl's terminology, occluded features are experienced as 'co-present' [1999 p. 222], because objects are given in perceptual experience as exhibiting 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*. (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'.)

Sensorimotor enactivists share Husserl's enthusiasm for this phenomenon. Noë's rather confusing term for it in *Action in Perception* is 'virtual presence', but I shall adhere to the original Husserlian terminology throughout. According to Noë's enactive theory of vision, the phenomenon of co-presence owes to the world's being experienced as 'available to perception through appropriate movement' [Noë 2012 p.58, italics removed], and this requires that perceivers possess sensorimotor understanding (also variably referred to as 'sensorimotor knowledge' or 'sensorimotor skill'). Being a perceiver is said to require an implicit grasp of sensorimotor contingencies - the law-like regularities 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]. This implicit grasp of sensorimotor contingencies is held to be a practical, as opposed to propositional, form of knowledge; in Ryle's idiom, it involves knowing-how rather than knowing-that [Noë 2004 pp.117-22, Ryle 2000 Ch. II]. On the enactive account, although we do not need to continually move in order to experience co-presence [Noë 2010], we experience co-presence because we know implicitly how to maneuver ourselves in relation to the object in such a way as to bring the occluded side or feature into view.

Noë's characterisation of sensorimotor understanding as practical knowledge of possibilities for action leads him to cash out the horizonal structure of visual experience in terms of Gibsonian affordances. For Gibson, to see an affordance is to directly perceive a familiar object's practical 'value' or 'meaning' [1986 p.127], that is, to see it as suggesting a possible usage which can be taken up in action: a chair affords sitting to a creature capable of sitting, lateral terrain affords walking to a creature capable of perambulation, and so forth. In Gibson's ecological theory of vision, though affordances are 'external' properties of objects, 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 enactive 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 their sort of approach to vision from more traditional cognitivist theorising is an emphasis on self-initiated action. For example, O'Regan is quick to defend his work on vision 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 [2010; 2012] now presents his work on perceptual 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 to 'a kind of dance' [2012 p.130]. In any case, the invocation of Gibson's ecological approach to 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, then, I will take it as an undeniable phenomenological fact that object 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 the horizonal structure of visual experience can be understood in terms of Gibsonian affordances and the possession of practical know-how. In the next section I outline three related objections to this claim. In the following section, I appeal to Merleau-Ponty's phenomenology to pick up the pieces.

3 Object Horizons Are Not Affordances

There are at least three 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, and not all bodily movement that would bring co-presented features to visual presentation requires volition. 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 voluntary. 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 their 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 would modify one's current perceptual experience need not require seeing an object as affording possibilities for active self-movement.

It is important to distinguish between 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's knowing how to act in certain ways, that is, if it contains some reference to possible self-initiated movement. Affordances are clearly action-dependent in this sense, for taking up an affordance requires an exercise of volition. 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 conceptual space between affordances –which presuppose volition, and object horizons, which do not, which gives us good grounds for refraining from equating the two. 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 movements of one's body and nearby objects. The crucial point is that this does not by itself license Noë's much stronger claim that object horizons are affordances.

A second problem with equating horizons with affordances concerns the psychological development of vision and agency. There are empirical grounds for holding that horizons are developmentally prior to affordances. Consider Held and Hein's [1963] famous 'kitten carousel' experiment. In contrast to a popular misconception the experiment concerned visually guided behaviour, not visual experience. For three hours daily ten pairs of neonatal kittens were placed in apparatus resembling a fairground carousal: a circular box with a two-pronged rotating arm fixed 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 surface, and displayed avoidance behaviour to looming objects, while the passive kittens did not, though their responses did normalise within forty-eight hours. Unsurprisingly, they conclude that 'self-produced movement with its occurrent visual feedback is necessary for the development of visually guided behaviour' [Held & Hein 1963 p.875].

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 [2004 p. 234 §9]. 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 [1995 pp.215-6] is that the cliff looks the same to both kittens, but only the active kittens have developed the association between the appearance of a flat surface and the feel of solid terrain under their paws.

Continuing to grant that object horizons are a basic, irreducible feature of visual experience, engaging in a spot of feline 'hetrophenomenology' (see Dennett [1991] pp.72-85] 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 horizons are phenomenologically more basic than 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. 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 horizons with affordances stems from an inconsistency in the attribution of practical knowledge to subjects incapable of performing the required bodily movements themselves. While Noë does not 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]

Noë frequently describes 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, as he argued that skills are acquired dispositions to act [2000 p.33]. For example, if I cannot play the guitar to a certain standard, then I do not possess the skill of guitar playing: I do not know *how* to play the guitar in the required sense. If 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. Skills are, in Merleau-Ponty's words, 'knowledge in the hands' [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 age has deprived her of her skill. As the following quotation makes clear, this Rylean conception of practical knowledge is explicitly endorsed by Noë[1]

> 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]

---

[1] Ryle's knowing-how/knowing-that distinction has been 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 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 suggests such an interpretation when he writes that *The Concept of Mind* as 'could be described as a sustained essay in phenomenology, if you are at home with that label' (2009 p.196).

The problem for Noë's sensorimotor enactivist should now be obvious: if the possession of practical knowledge is 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.

Locked-in syndrome is instructive in this regard. The 'classical' variation of locked-in syndrome involves complete paralysis apart from blinking and limited vertical eye movement [Bauer, Gersenbrand and Rumpl 1979].[2] Sufferers of the syndrome can communicate using systems of blinks and vertical eye movements and with the help of various eye-tracking technologies [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 enactivists need not, do not, and given their phenomenological starting point cannot, deny this. But it is wrongheaded 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 maneuver their body around an object in the sense of having the required skills; were they miraculously cured they would need to reacquire them through practice and physiotherapy. Practical knowledge 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. 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 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. Their visual experience is not that of a human equivalent of Hein and Held's passive kittens. 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 practical

[2] 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).

know-how. What is needed is a better philosophical framework in which to make sense of sensorimotor understanding.

4 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 practical 'value' which the objects of vision have for the perceiver who possesses and exercises the relevant practical know-how – affordances – persists even after this know-how has been lost. The enactivist project is therefore hampered by Noë's restricted conceptual toolkit. *Phenomenology of Perception* gives us the additional tools needed to untangle the knots in which Noë's sensorimotor enactivist ties herself by attempting to conceive of sensorimotor understanding solely in terms of Rylean practical know-how and Gibsonian ecological psychology. This requires some preliminary exposition on the aim of phenomenology considered not as a subject matter, but as a discipline.

Phenomenological philosophy is personal-level analysis par excellence, but there is considerably more to it than introspective reports on the content or character of psychological states. Husserl's philosophical project gradually evolved from 'static' into 'genetic' phenomenology, and Merleau-Ponty's *Phenomenology of Perception* sits squarely in the latter category. Static phenomenology describes experience in terms of hypostatised appearances. Genetic phenomenology goes further, and attempts 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 limited to the description of its end products. Interestingly, Noë, who takes himself to be 'investigating the phenomenology of perceptual experience' [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 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 body schema and sedimentation.

In contrast to a percept or mental representation of one's own body, the body *schema* is Merleau-Ponty's term for the integrated system of pre-reflective bodily capacities that structure perceptual experience. The notion of a body schema has received considerable attention in contemporary cognitive science thanks to the work of Shaun Gallagher who, following Merleau-Ponty's lead, champions a distinction between body schema and body image.[3] Though Merleau-Ponty would certainly agree with Gallagher that the body schema should be distinguished from any mental state which has the body as its object, the body schema has for Merleau-Ponty an existential 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' [Gallagher 2005 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' [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 specific 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].

Crucially, the body schema is adaptable. By honing skills and acquiring new habits, it can be 'reworked and renewed' [ibid. p.143], and this endows the perceiver's 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 the idea of 'bodily meaning', an example of the body schema's contribution to lived experience should help to clarify what Merleau-Ponty means by this. For expert musicians, perceptual encounters 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 shapes, picking techniques and scale patterns lacking in non-musicians who are nevertheless perfectly capable of performing similar finger movements. Consequently, guitars are perceived in a more meaningful way by the guitarist – they draw on a richer sensorimotor understanding alien to a perceiver lacking this enriched motor signification. The non-guitarist knows (in the propositional sense of 'knows' – knowing-that) very well what the guitar is for, and might even know something of how it is played, but the guitar does not afford playing for them in the concrete sense experienced by the skilled player for whom the guitar represents a genuinely possible motor project. The same goes, mutatis mutandis, for other forms of skilful sensorimotor interactions. As a non-driver, the interior of a car is mysterious to me in a way which invites laughter from my road-ready 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

[3] '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)

recall similar experiences where one struggles to 'interpret' one's environment while others negotiate it effortlessly (the reader is invited to think of their own examples). 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 engagement. It is by virtue of the body schema and its adaptability that we are geared into our environment in such a way that it *makes 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' [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 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. [2012 p.131]

Merleau-Ponty's discussion of language provides a handy illustration of sedimentation at work [see 2012 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 these 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*), 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 nihilo, 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 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.

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, 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 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* – although the opportunity to form new motor significations is 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*. The second comes from his discussion of 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, than ensures that we are situated within all these relationships. [ibid. 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, 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 experience of interactions 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 a '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 untenably impoverished account of the phenomena.

5 Conclusion

I have argued that Noë's account of sensorimotor understanding suffers from deficiencies and inconstancies which Merleau-Ponty's early phenomenology is equipped to rectify, albeit at the expense of incurring an inflated conceptual inventory which potentially carries its own distinct set of philosophical problems. While Noë is certainly correct that 'the task of phenomenology ought to be…to catch experience in the act of making the world available' [2004 p.176], this requires explicating not just how perceivers 'bring the world forth' [2012 p.14] by applying their sensorimotor understanding in experience, but also what it is to be embodied and situated in such a way as to make such understanding possible in the first place, and to sustain the meaningful structure of perceptual experience to which it give rise. These questions, though not necessarily beyond the scope of cognitive science, are fundamentally existential ones, and the beginnings of answers to them are only sketched here. The sensorimotor theorist, who, like Noë, aspires also to be a phenomenologist, has their work cut out for them.

References

1. Baldwin, T.: Speaking and Spoken Speech. In: Baldwin, T. (ed.) Reading Merleau-Ponty, pp. 87–104. Routledge, London (2007)
2. Bauer, G., Gerstenbrand, F., Rumpl, E.: Varieties of the Locked-In Syndrome. Journal of Neurology 221(2), 77–91 (1979)

3. Chisholm, N., Gillett, G.: The Patient's Journey: Living With Locked-In Syndrome. British Medical Journal 331, 94–97 (2005)
4. Dennett, D.: Consciousness Explained. Penguin, London (1991)
5. Gallagher, S.: How the Body Shapes the Mind. Oxford University Press, Oxford (2005)
6. Gibson, J.: The Theory of Affordances. In: The Ecological Approach to Visual Perception, pp. 127–143. Psychology Press, NJ (1986)
7. Held, R., Hein, A.: Movement-Produced Stimulation in the Development of Visually Guided Behaviour. Journal of Comparative and Physiological Psychology 56(5), 872–876 (1963)
8. Husserl, E.: Cartesian Meditations. Trans. D Cairns. Kluwer, Dordrecht (1988)
9. Husserl, E.: Horizons and the Genesis of Perception. Trans. A. Steinbock. In: Welton, D. (ed.) The Essential Husserl, pp. 221–233. Indiana University Press, IN (1999)
10. 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., Goldman, S.: The locked-in syndrome: what is it like to be conscious but paralysed and voiceless? Progress in Brain Research 150, 495–511 (2005)
11. Merleau-Ponty, M.: The Structure of Behaviour. Trans. A. Fisher. Methuen, London (1965)
12. Merleau-Ponty, M.: Phenomenology of Perception. Trans. D. Landes. Routledge, London (1945, 2012)
13. Noë, A.: Varieties of Presence. Harvard University Press, MA (2012)
14. Noë, A.: Vision without Representation. In: Gangopadhyay, N., Madary, M., Spicer, F. (eds.) Perception, Action, and Consciousness, pp. 245–256. Oxford University Press, Oxford (2010)
15. Noë, A.: Action in Perception. MIT Press, MA (2004)
16. O'Regan, J.: Explaining What People Say About Qualia. In: Gangopadhyay, Madary, and Spicer, pp. 31–50 (2010)
17. O'Regan, J., Noë, A.: A Sensorimotor Account of Vision and Visual Consciousness. Behavioural and Brain Sciences 24, 939–1031 (2001)
18. Patterson, J., Grabois, M.: Locked-in syndrome: a review of 139 cases. Stroke 17, 758–764 (1986)
19. Ryle, G.: Phenomenology versus 'The Concept of Mind'. In: Critical Essays: Collected Papers, vol. I, pp. 186–204. Routledge, London (2009)
20. The Concept of Mind. Penguin, London (1949, 2000)
21. Smith, E., Delargy, M.: Locked-in syndrome. British Medical Journal 330, 406–409 (2005)
22. Stanley, J., Williamson, T.: Knowing How. Journal of Philosophy 98(8), 411–444 (2001)

How Enactive Is the Dynamic Sensorimotor Account of Raw Feel?: Discussing Some Insights from Phenomenology and the Cognitive Sciences

Alfonsina Scarinzi

Georg–August Universität Göttingen, Germany
alfonsinascarinzi@googlemail.com

Abstract. This contribution deals with the question of how enactive O'Regan's dynamic sensorimotor account of phenomenal consciousness is. It answers this question by focusing on O'Regan's dynamic sensorimotor account of raw feel. It supports the view that O'Regan's sensorimotor approach is *semi-enactive* because it clings to environment-centric ecological Gibsonian overtones. It emphasizes the instrumental role of action in perception enactivism rejects and neglects the subjectivity of experience. This contribution makes the point that the role of the motor and cognitive-emotional aroused lived body and the subject's conscious access to it in experiencing the qualities of sensorimotor interactions and hence the subjectivity of experience need to be taken into account in order to overcome the ecological environment-centric overtones of O'Regan's approach.

Keywords: enactivism, ecological approach, sensorimotor approach, lived body, phenomenology, enactive approach to emotions, anti-dualism.

1 Introduction

It has become a commonplace in research in cognitive science and philosophy of mind to use the term 'enact' or 'enactive' to refer to intrinsically active perception and to the understanding of cognition as based on knowing how and hence on understanding what enables us to move and to engage with the world we co-determine through our sensorimotor skills and abilities [10] [8]. As Hutto [10] remarks, 'enactive' means that we know how to tie our shoes, to ride a bike, to play table-tennis without following propositional rules based on inner representation of knowledge about the world. More precisely, 'enactive' refers to a framework within cognitive science called non-classical cognitive science or enactivism, whose theoretical tenets Steve Torrance [22] describes in the following way: "(a) Minds are the possessions of embodied biological organisms viewed as autonomous self-generating and self-maintaining – agents. (b) In sufficiently complex organisms, these agents possess nervous systems working as organizationally closed networks, generating meaning, rather than processing information as inner representations of the external world.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_5,

(c) Cognition, conceived fundamentally as meaning-generation, arises from the sensorimotor coupling between organism and environment. (d) The organism's world is 'enacted' or 'brought forth' by that organism's sensorimotor activity; with world and organism mutually codetermining one another, in ways that have been analyzed by investigators in the continental phenomenology tradition. (e) The organism's experiential awareness of its self and its world is a central feature of its lived embodiment in the world, and therefore of any science of the mind."

Enactivism conciliates phenomenology and cognitive science acknowledging that especially the phenomenological studies on the lived body[1] can clarify and guide scientific research on subjectivity and consciousness [19] [23].

In his last work with the title *Why Red Doesn't Sound Like a Bell*? Kevin O'Regan[2] wonders if he can be considered an enactivist, for in his sensorimotor account of phenomenal consciousness he acknowledges the central role of action in perception enactivists support, but also the usefulness of representations enactivists reject. How enactive is O'Regan's dynamic sensorimotor account of phenomenal consciousness?

The aim of this contribution is to answer this question by focusing on O'Regan's dynamic sensorimotor account of raw feel. According to O'Regan [16], "raw feel" is whatever people are referring to when they talk about the most basic aspects of their experience." (96). It corresponds to the "something it's like"–sensation. An example is the feel of red, this is what is at the core of what happens when I look at a red patch of color beyond cognitive states and bodily reactions. In O'Regan's work raw feel corresponds to qualia.

In this contribution I will make the point that O'Regan's sensorimotor approach in general and sensorimotor account of raw feel in particular are *semi-enactive*. I will argue that they cling to environment-centric ecological overtones emphasizing the instrumental role of action in perception enactivism rejects and are hence incomplete. In my view O'Regan's approach needs to take into account the role of the motor and cognitive-emotional aroused lived body and the subject's conscious access to it in experiencing the qualities of sensorimotor interactions in order to overcome its ecological overtones, which are not enactive.

In the following I will first analyze O'Regan's dynamic sensorimotor account of raw feel [16] in the light of enactivism taking into account what enactivism says about mental representation and embodiment in order to point to the main differences between the enactive approach and O'Regan's sensorimotor approach. I will then turn to O'Regan's [16] *phenomenality plot* – the graph indicating qualities of experience which are objectively quantifiable by physical and physiological measurements – and analyze his sensorimotor approach to emotions as an example of raw feel of experience in the light of the enactive research on the phenomenological lived body and its role in the embodied enactive approach to emotions. I will point to the fact that the main consequence of bringing enactivism and the sensorimotor approach to raw

---

[1] See footnote 5.

[2] O'Regan [16] (p. 68 note 1); O'Regan together with Alva Noë developed the well-known sensorimotor account of consciousness bridging the explanatory gap consisting in the problem of explaining qualia in physical or biological terms. See [15]

feel closer to each other consists in leading to the development of a method based on a combination of third-person and first- person approach for the investigation of the feels of experience.

2 The Dynamic Sensorimotor Account of Raw Feel as *Semi-enactive*

The theoretical background of the dynamic sensorimotor account of raw feel is the sensorimotor contingency theory of perceptual experience or sensorimotor account of consciousness by O'Regan and Noë [15]. This is the view according to which perceptual experiences are ways of acting, constituted in part by the perceiver's skillful mastery of the relation between sensory experience and movement. The senses have different characteristic patterns of sensorimotor dependence, and perceivers have an implicit, skillful mastery of these differences. For each modality of perceptual experience – seeing, hearing, touching, and so on – there is a corresponding pattern of sensorimotor interdependence that is constitutive of that modality. What it is to experience the world perceptually is to exercise one's bodily mastery or know-how of certain patterns of sensorimotor dependence between one's sensing and moving body and the environment [15] [16].

In the same way as perception in this approach is considered to be something we do, O'Regan [16] considers also feel and raw feel as something we do. Experiencing raw feel involves engaging with the world (67; 113). Sensory inputs feel like something because the so called attributes of sensorimotor interaction we have with the environment provide the quality of sensory phenomenality to sensory inputs. O'Regan [16] identifies four attributes of the sensorimotor interaction which are specific to what we call feel and occur together when we experience a raw feel. These are richness, bodiliness, (partial) insubordinateness and grabbiness. Richness refers to the variety of the qualities of real-world sensory interactions. Bodiliness expresses the dependence between body motion and sensory input in a sensory modality. It is a distinguishing feature of neural activity deriving from external-world stimulation. (Partial) insubordinateness refers to the fact that real-world interactions are subordinate to our voluntary body motions, but they can cause changes in sensory input without us moving. Grabbiness is the capacity of a sensory modality of grabbing our cognitive processing. It is the tendency of something to attract one's attention. These qualities of interaction determine the sensory presence and hence the raw feel of an experience. According to O'Regan [16] such qualities are the physical qualities of an interaction and are hence objective and measurable (176), which gives the sensorimotor approach its advantage over an approach based on the neural correlates of consciousness. Even if in this contribution I will not compare the sensorimotor approach with an approach based on the neural correlates of consciousness, I shall come back to what I consider to be a methodological incompleteness of the objective measurement of the qualities of experience, which neglects the first-personal subjective character of experiencing raw feel, later on.

Both Thompson [20] and Hutto [10] observe that the sensorimotor approach is not properly enactive or enactivist. Thompson observes that even if it allows to explain the enactive subject's sensorimotor coupling with the world in terms of skillful mastery of the relation between sensory experience and movement, it is incomplete because it lacks a notion of an experiencing agent and it lacks to take into account the pre-reflective nature of bodily self-consciousness. Hutto points to the fact that in spite of the central role of skill–based explanations, the sensorimotor approach is riddled with suppositions threatening to reduce it to a rules-and-representations approach, which is not in line with antirepresentationalism enactivism supports. My aim in this section is not to show how the sensorimotor approach can be transformed into a properly enactive one. Rather, taking into account Thompson's and Hutto's observations as a starting-point, in this section I will answer the question of how enactive with respect to the role of mental representation and embodiment in enactivism the sensorimotor approach and the sensorimotor approach to raw feel are.

My claim is that they are *semi-enactive* because they are essentially limited in scope. In order to be able to bring them closer to enactivism sensorimotor research needs to take into account the motor and cognitive-emotional role of the lived body and the subjective access to it in the investigation of the agent's or subject's experience of the qualities of interactions. In other words, developing an account of subjectivity would make them enactive.

2.1 The Role of Representation

What is the role of mental representations in O'Regan's sensorimotor approach and in his sensorimotor approach to raw feel? In the following I will argue that in O'Regan's approach its use is superfluous and its role weak and that the weak role of representation contributes to determining the ecological Gibsonian and *semi-enactive* nature of O'Regan's sensorimotor approach to raw feel.

O'Regan [16] rejects the copy view of experience and hence the strong view of representation. This is the idea that experience is somehow constituted by the formation of passive, internal representations of outer scenes and of a pre-given world. Rejecting this idea is in line with enactivism, which does not follow the idea of assuming that the world is pre-given, that its features can be specified prior to any cognitive activity hypothezising the existence of mental representations inside the cognitive system to explain the relation between this cognitive activity and a pre-given world [23] [8]. The strong view of representation presupposes a dualistic discontinuity between the body and the mind [12] [9]. But the strong idea of the existence of mental representations about the outside world is considered to be implausible. As a matter of fact, empirical research in neuroscience, situated robotics, ecological psychology, developmental psychology, philosophy of mind suggests that there is no single center of thinking. Kirchhoff [12] formulates the two main arguments coming from empirical evidences against representationalism in the following way: "The first of these arguments, the threat from *nontrivial causal spread*, occurs whenever the material vehicles of cognitive architecture are *causally spread* beyond the brain and *nontrivially* involved in the completion of cognitive tasks. The second of these arguments, the threat from *continuous reciprocal causation*, occurs whenever the causal contributions made by components of a system partially determines and is partially determined by causal contributions of other systemic components, thereby making it impossible to assign a specific subtask to an identifiable subsystem within a larger system".

In other words, in more philosophical terms we can say that thoughts, propositions, concepts and percepts can not be considered to be quasi-objects about an external pre-given world. Rather, they are in and of the world as patterns of experiential interactions of organism-environment couplings[3] that constitute experience [11]. This can be considered to be the antirepresentational basis of enactive cognitive science Varela et al. [23] formulate in their work *The Embodied Mind*, according to which a cognitive system is understood on the basis of its so called "operational closure" which means that the results of its processes are those processes themselves. This means that a cognitive system does not operate by representations. Rather, it enacts or brings forth a world as a domain of distinctions according to its bodily structure. In other words, a cognitive system creates a minimal distinction between an interior and an exterior, and guarantees the continuous dynamical generation of its stable "internal coherence" [17] in the relation of co- determination with the environment.

Even if O'Regan [16] rejects the copy view of experience and hence the strong notion of representation, he clings to the view that representation is useful depending on how it is used. O'Regan [16] seems to support the view that the use of representation should be restricted to indicate the patterns of a structure and not the copy of a pre-given world. This is also called the weak view of representation. This can be considered to be both in line with enactivism, with which the weak notion of representation is compatible, and in line with the use of representation in the landscape of naturalistic contemporary theories of cognition and perception [11]. Nevertheless, O'Regan [16] comes to the conclusion that it is better to avoid words like representation, which give rise to confusion. Hence, he considers the use of representation as superfluous and prefers focusing on better explaining what is meant by 'having a sensory experience'. He writes: "[...] what we really mean is just that we are currently involved in extracting information from the environment in a way which is peculiar to the visual modality." (64)

The important point here for my claim that O'Regan's approach remains *semi-enactive* despite the acknowledgment of the weak notion of representation is the expression "extracting information from the environment". In O'Regan's approach extracting information from the environment is possible thanks to action, to bodiliness and hence to movement and its sensory consequences. In spite of the central role of action and bodiliness, as extracting information from the environment is described as a perceptually guided action, the environment seems to be independent from the perceiver or perceiving subject. Information is there to be discovered and movement is the instrument to do this. As a matter of fact, with reference to the sensory modality of seeing, O'Regan [16] writes: "If I move (but I needn't actually do so), I (implicitly) know there will be a large change in my visual input." (67) The instrumental role of

[3] See footnote 4.

action in the sensorimtor approach to perception is explained in a clear manner also by O'Regan's co-author Alva Noë [14] in his work *Action in Perception*, where he points out that our sense of the perceptual presence of a cat, for example, requires that we take ourselves to have access to the whole cat and the ground of this accessibility is our possession of sensorimotor skills (63).

It is necessary to remark here that the instrumental role of movement cannot be considered to be ‘enactive’ [13], because it implies that information cannot be brought forth by the perceiver in a relation of co-determination with the environment. It can just be picked up and hence extracted from the environment. Interaction through movement having an instrumental role is just a sort of direct detection, which reflects the Gibsonian ecological approach.

The instrumental role action and movement have in O'Regan's approach leads me to believe that O'Regan's approach is closer to ecological Gibsonian approaches to perception than to enactivism. As Chemero [1] makes clear, according to the Gibsonian direct approach to the perceptual detection of information, the environment contains sufficient information to guide the subject's behavior and no information is added in the mind (106). Mental representation in the strong sense is hence not necessary, while representation in the weak sense is compatible with direct perception the sensorimotor approach supports. I would like to remark here that in spite of the acknowledgement of perception as a direct process, ecological approaches are not enactive. Despite the tendency in research to draw analogies between the ecological and the enactive approach, Varela et al. [23] consider their enactive approach not to be ecological at all (203). They briefly summarize that in their enactive approach the environment is not independent. Rather, it is enacted by histories of coupling.[4] This is the main relevant difference. Varela et al. [23] claim that the environment is sensorimotor enactment. They point out that they build up the theory of perception from the structural coupling of the animal by specifying the sensorimotor patterns that enable action to be perceptually guided. This means that movement has a constitutive role in the determination of the perceived environment and hence of what O'Regan refers to as information [13]. The enactive approach is a relation-centric approach, a middle-way between the environment and the perceiver [8].

From my point of view, the sensorimotor approach clearly conveys a *semi-enactive* idea, for it rejects the strong version of representation, integrates or considers to be possible a weak version of representation, but the instrumental role of perceptually guided action to extract information from the environment does not reject the view of an environment-centric sensorimotor approach to experience.

In other words, with reference to the sensorimotor approach to raw feels, perceiving the raw feel of red is something we do by extracting information from the environment according to a process of direct perception in the ecological sense.

I am making the point here that in the sensorimotor approach action and movement just give the quality of experience in a pre-given world its sensory presence in the

[4] In the enactive approach sensorimotor coupling refers to the type of interaction with the environment by which agents actively generate their identity by selecting from the environment their viable world called “cognitive domain” that is brought forth or enacted by that agent's autonomous mode of coupling with the environment. See [18] [23].

environment. More precisely, they allow the subject to gather it from the instrumental action-related changes in the sensory input. This is not enactive despite the antirepresentational overtones of the sensorimotor approach.

2.2 The Embodied Mind and the Sensorimotor Approach to Raw Feel

In order to be enactive the sensorimotor approach to raw feel needs to be based on the notion of embodied mind [23]. Does O'Regan [16] take the embodied mind into account? In the following I will claim that O'Regan's approach is compatible with the embodied mind thesis, but it does not take it into account because it remains an approach which, unlike enactivism, does not aim at solving the problem of the Cartesian mind-body dualism. Rather, it focuses on providing a framework in which problems appear as non-problems [10], which is far away from being a specific alternative framework to dualistic approaches.

In enactivism the embodied mind thesis is traced back to the work *The Embodied Mind* by Varela et al. [23]. It refers to the claim that perception, thinking, feelings and desires – that is the way we behave, experience and live the world – are contextualized by our being active agents with the particular kind of body we have. It rejects the Cartesian mind-body dualism according to which mind and body are split into segregated, pure forms. Putting the mind into the body means that in the interaction with the environment the engaged human living body is inconceivable without a mind. According to the embodied mind thesis the body is a form of lived experience, references a biological standpoint [25] and as lived experience also a phenomenological or psychological person standpoint [6] [20]. As Fuchs [6] points out, the body is not simply the carrier of the brain (163). Rather, it is organized in such a way that it displays the suitable structures to produce the conscious manifestation of life. It is useful to remark here that the embodied mind is relational, distributed over body, brain and environment, does not reference merely physical structures, is not about a disengaged agent defined by its movements. As Thompson [19] remarks, in the embodied enactive approach to the mind the inner and the outer are not separate spheres but mutually specifying domains enacted by the structural coupling of the system and its environment. Cognition is embodied action. In other words, the lived body (the inner) is a dynamic condition and a performance of the living body in the interaction with the environment (the outer) in a relation of co-determination. As Fuchs [6] points out, subjectivity is necessarily embodied, so living body is necessary subjective (163). The embodied mind is hence not simply about a moving agent peeking at a preformed world and drawing meaning directly from that world. It is not simply direct perception determined by the instrumental action-related changes in the sensory input.

In considering raw feels simply as qualities of experience constituted by skilled modes of interaction with the environment O'Regan [16] is mainly concerned with finding a way to avoid reduce the explanation of the qualities of experience to brain mechanisms and the role and existence of neural circuitry and special neurons and hence in overcoming or bridging the explanatory gap. Nothing more. He writes: "But the quality of the feel involved is not caused by the activity of the brain mechanism; it is constituted by the quality of the interaction that is taking place and that the brain mechanism has enabled" (114). In my view, also in stressing that the qualities are objective he focuses only on the fact that they are objectively measurable by a physicist. I believe that in rejecting the strong version of representation and following the direct approach to perception the sensorimotor approach rejects de facto the mind-body dualism, but in considering the objective instrumental action-related changes in the sensory input only without taking the subjectively lived body into account the sensorimotor approach to the qualities of experience remains partial.

In other words, this partially embodied sensorimotor approach to raw feels lacks an explanation of how the experience of sensorial events relate one's subjectively lived body[5] to itself and hence how the subject makes sense of the objective qualities of his/her own experience. In my view, this is surprising because experience is always someone's experience. O'Regan's approach completely neglects the subjectively lived body's unique status as a physical subject. I believe that what is missing in this approach, is something immanent to the system, shaping its way of being in the world, its way of being coupled. In order to be enactive, O'Regan's approach needs to take into account the realm of consciousness. This is the part of our cognition that we access from a subjective point of view [17]. The subject experiencing the qualities of interactions (the experiencing agent) does this from a subjective embodied point of view showing also certain bodily reactions to an external observer, which are objectively measurable markers. The careful examination of the subject's or agent's experience from the point of view of an external observer requires hence to take into account a first-person method of verbal explicitation of experience, like the report of the observation of one's own lived experience while experiencing the raw feel of the qualities of interaction.[6] Moreover, it also requires a combination of first-person and third-person methodologies from the point of view of an external observer. The reasons are described by Varela and Shear [24] as follows: “What we take to be objective is what can be turned from individual accounts into a body of regulated knowledge. This body of knowledge is inescapably in part subjective, since it depends on individual observation and experience, and partly objective, since it is constrained and regulated by the empirical, natural phenomena.” [24]

Summing up, one can say that combining subjectivity and objectivity is not O'Regan's ambition. Hence, also from the point of view of embodiment O'Regan's approach is *semi-enactive*.

---

[5] The lived body is also called *Leib* in phenomenology. It is your own body as experienced by yourself.

[6] In their work with the title The validity of First-Person Descriptions as Authenticity and Coherence, Journal of Consciousness Studies, 16, No. 10–12, pp. 363–404 (2009) Claire Petitmengin and Michel Bitbol analyze the reliability of first-person descriptions. They come to the conclusion that their validity can be measured in dynamic terms of performative consistency of the acts which produce first-person research. They write: “[…], researchers in the domain of lived experience cannot avoid making a detour by their own experience. Their expertise must not limit itself to the inventory of objective signs, but must extend to the exploration of their own subjectivity.”

3 The Phenomenality Plot as *Semi-enactive* and the Case of 'Enactive' Emotions

In order to better illustrate the *semi-enactive* nature of the sensorimotor approach to raw feel and the relevant role of the lived body O'Regan [16] completely neglects, in this section I will put into focus the limits of O'Regan's *phenomenality plot* from an enactive viewpoint and his way of applying the sensorimotor approach to raw feel to emotions. I will illustrate this using the enactive approach to emotion by Colombetti [4] with the purpose of discussing how the notion of lived body can help better explain the constitutive role of body motion in the feel of experience.

In his work O'Regan takes emotion into account because it is a special case, for its feel is not phenomenally present and not precisely localized. This remark by O'Regan will be my starting-point to claim that without the notion of the lived body the sensorimotor approach to raw feel remains incomplete. My claim is in line with Thompson's [20] remarks about the sensorimotor approach to consciousness mentioned above.

3.1 The Bodiliness and Grabbiness of an Emotion in the *Phenomenality Plot*

The fact that the feel of the experience of emotion is not phenomenally present and not precisely localized is a good condition to explore both the central and the partial and insufficient role of bodiliness and grabbiness in explaining the feel of experience. It is also a good starting-point to discuss why they are necessary but not sufficient conditions to explain the feel of experience, which has an objective but also a subjective complementary aspect.

O'Regan [16] considers richness, bodiliness, insubordinateness, and grabbiness as the basis for the "what it is like" of sensory feels. The critical factors remain bodiliness and grabbiness, which are diagnostic of the degree of sensory presence of the different experience, and which he plots on a graph he calls the *phenomanality plot* where he indicates the amount of bodiliness of a given type of mental or neural activity and the amount of grabbiness that activity has. In the paragraph about emotion in his work O'Regan focuses basically on the feeling of an emotion, which is considered to involve cognitively monitoring of the bodily reactions associated with an emotion. The only ambition O'Regan has is to show that the sensorimotor approach to the feel of emotion can explain the cognitively monitoring of the bodily reactions associated with an emotion. It can explain that the experienced feel of an emotion has to be searched in the ongoing interaction with the environment whose qualities constitute the experienced feel. What does this mean? And how enactive is it?

If we consider "fear" as a prototypical emotion, as O'Regan [16] in his work does, bodiliness and grabbiness are the two elements which we have to use to find out how the experienced feel of "fear" comes into being. According to O'Regan, a subject can recognize "fear" because a situation attracting her attention is interpreted as dangerous. It is hence a social interpretation of interaction also based on previous experiences. Bodiliness plays a secondary role, as moving the body does not change the subject's fear, unless movement is instrumental to moving the perceiving body out of the dangerous situation. In this case movement would change the conditions for interaction. In the sensorimotor approach hence the cognitively monitoring of the bodily reactions associated with an emotion tends to reduce the investigation of the feel of an emotion to the objective qualities of interaction an experiencing agent has with the environment and which are measurable from a third-person point of view. In other words, according to the sensorimotor approach, in order to have the required sensitivity for the arousing emotion the precondition the subject needs is the abilities such as skills to engage in emotion-laden interactions. From my point of view, in O'Regan's approach bodiliness and grabbiness can contribute to monitoring bodily reactions associated with an emotion in a quite weak way, for they are not able to explain the relation between the moving body in the environment and the evaluation of the situation the subject interacts with.[7] In other words, I believe that if one can not find a way to explain how it is possible that the experiencing agent or subject evaluates and hence recognizes the situation she interacts with as fear provoking or joy provoking or anxiety provoking, the sensorimotor approach remains partial from an enactive point of view, for it does not explain how the outer – the environment – and the inner – the subjectively lived body – determining embodied meaning and sense-making in interaction mutually specify each other in interpreting an emotionally charged situation. The case of the feel of emotion is a good example for the limits of the sensorimotor approach. In my view, O'Regan's approach remains environment-centric without being able to localize the feel of emotion and its sensory presence in the interaction with the environment. The neglected point is that the feel of emotion is someone's feel. What is needed is a way to take this subjectivity into account. The phenomenological notion of the lived body and the enactive cognitive-emotional access to it allow to take such a subjectivity into account without excluding the simultaneous role of bodiliness and grabbiness in the interaction with the environment.

In other words, I believe that the sensorimotor approach to the feel of experience and hence also to the experienced feel of an emotion can be brought closed to enactivism if movement can be considered to be constitutive and not instrumental in the determination of the feel of the qualities of interaction and hence of experience. In the following I will consider the case of the enactive approach to emotion with the purpose of showing that the role of the lived body is essential to explain the co-determination of the inner and the outer, the subjectivity of the feel of experience and the constitutive role of motion in such a relation of co- determination.

3.2 The Cognitive-Emotional Role of the Lived Body: Towards the 'Enactive' in the Sensorimotor Approach to the Feel of Experience

Giovanna Colombetti and Evan Thompson [3] summarize with the following words the main claims of enactive research in the field of emotion theory: "Cognition is a form of embodied action [...]. The enactive approach implies that we need to move

[7] According to Fingerhut [5], the problem with the *phenomenality plot* is its rigidity, which leads to a weak explanation of the correlation between presence and bodily modulation (178).

beyond the head/body and subjective/objective dichotomies that characterize much of emotion theory. Appraisal is not a cognitive process of subjective evaluation "in the head" and arousal and behavior are not objective bodily concomitants of emotion. Rather, bodily events are constitutive of appraisal, both structurally and phenomenologically." (56–58)

Enactive research on cognition, emotion and the lived body has emphasized that emotion and cognition or also bodily sense-making[8] and emotions are embodied and interdependent. Thompson [19] in explaining the approach by Marc Lewis formulates this idea in the following way: "Cognitive and emotional processes modify each other continuously on a fast timescale, while simultaneously being constrained by the global form produced by their coupling in a process of circular causality. This emergent form, the emotional interpretation, is a global state of emotion-cognition coherence, comprising an appraisal of a situation, an affective tone, and an action plan." (371)

According to Colombetti's enactive approach to emotion and bodily sense- making [4] [18], bodily sense–making, which is central in the constitution of an emotion and hence of degrees of value in the evaluation of a situation, manifests itself in the experience of the aroused lived body and hence in experience through embodied emotions such as fear, anger, happiness, guilt, anguish. These are the way the subject evaluates bodily sense-making in the interaction with a situation. The mentioned embodied emotions are bodily mediated cognitive-emotional evaluations of the bodily sense-making of an adaptation to environmental factors the organism interacts with in the environment and of their viability. They allow to subjectively feel the cognitive-emotional qualitative dimension of the degree of value of our interaction with different environmental factors through the aroused lived body. In developing the idea of the cognitive-emotional embodied evaluation, which the aroused lived body is a vehicle of, Giovanna Colombetti [4] [2] points out that as there is no cognition without emotion and emotion is embodied arousal needs no appraisal to be interpreted by the subject.[9] The aroused body is immediately available as such to the subject's experience through the evaluation of the bodily aspects of emotion as part of the subject's evaluation of the experienced world. Bodily arousal subsumes the whole subject's organism capacity to make sense of her world and is possible thanks to the lived body. In other words, as Todres and Galvin [21] following Eugène Gendlin

---

[8] In enactive research bodily sense-making refers to the process according to which the whole organism is a vehicle of meaning which is dynamically constructed by the subject having a perspective on the world. In the interaction with and adaptation to the environment bodily sense-making is the evaluation of an adaptation and takes place in the organism's coupling with the environment. It has both the function to contribute to maintaining the organismic integrity of the subject (regulation) and to expand the subject's cognitive domain through the active selection of viable environmental factors to be integrated into the subject's cognitive domain. [23] [18]

[9] According to traditional non-enactive dualistic emotion theories the notion of emotion is constituted simultaneously by a mental and a bodily event. The mental side of an emotion is called appraisal and the bodily side is called arousal. Without a cognitive activity there can be no emotion, there can be just bodily arousal. See [3]

point out, the conscious experience of the subjectively lived body is the subject's access to the pre-reflective dimension of experience. It is the access to the own multisensorial pattern of feeling that is the basis for different felt qualities or felt sense of a situation and for the emergence of thought and understanding.

Colombetti points out that in the same way as the pre-reflective lived body allows the experience of becoming aware of my body as that through which, for example, the experience of typing on the computer is possible, it allows to be similarly aware of the bodily arousal as that through which I am living the situation of an interview as anxiety provoking. As also Mark Johnson [11] points out, motor intentionality and the subjectively experience of the feeling of qualities of an experienced situation enable to account for meaning as grounded in bodily experience.

Against this background, we can say that in the process of the cognitive-emotional evaluation of and adaptation to the environment, the central role of the lived body is twofold:

1. it is the pre-reflective backdrop against which the perceptual and motor experience is constituted, which is its classical role, traditionally studied in phenomenology; the role of the lived body as a backdrop against which actions and experiences in the world take place becomes vividly apparent in the pre-reflective bodily self-consciousness: one's body shows itself to be a material thing animated from within by sensation and motility. Evan Thompson [19] gives the well-known example of a cup of hot tea. When I pick up a cup of hot tea, I feel the hot smooth surface of the porcelain and the heat penetrating my finger, and these sensations linger for a time after I have put the cup back down on the table. Such bodily experience is twofold: it is the experience of physical events that relate one's body to things and it is the experience of sensorial events that relate one's subjectively lived body to itself. Following Husserl the lived body manifests itself in perceptual experience as an implicit I can of movement and motor intentionality. Bodily subjectivity is "an I can and do move in such and such a way" (249).
2. it is the pre-reflective backdrop against which the cognitive-emotional evaluation of the experienced world takes place, which is its role in the enactive approach to emotion.

With reference to the subjectively felt qualitative dimension of experience, this means that this takes place against the backdrop of both a pre-reflective motor and cognitive-emotional lived body. Nevertheless, the cognitive-emotional evaluative pre-reflective dimension of the aroused lived body is the distinctive backdrop against which a viable evaluation and hence the cognitive-emotional qualitative dimension of the degrees of value of different environmental factors are subjectively felt in the experience and perception of the world, which are subordinated to motion.

What does this overview on the state-of-the-art in enactive research on emotion and cognition mean for my concern here? As I argued above, the role of the lived body is essential to explain the co-determination of the inner and the outer and the subjectivity of the feel of experience and that O'Regan's sensorimotor approach to the feel of experience and hence also to the experienced feel of an emotion can be brought closer to enactivism if movement can be considered to be constitutive and not instrumental in the determination of the feel of the qualities of interaction and hence of experience.

Combining the double role of the lived body as the pre-reflective backdrop of perceptual and motor experience and as the pre-reflective backdrop of the subjective cognitive-emotional evaluation of the experienced world with the sensorimotor role of bodiliness and grabbiness in interaction in the experience of the feel of an emotion and hence the inner with the outer can contribute to reaching this aim.

With these changed premises I would like to consider here again the experienced feel of "fear" and bodiliness and grabbiness. Bodiliness and grabbiness can be considered to be convincing markers of phenomenality in the enactive sense, where cognitive science and phenomenology are brought together to investigate consciousness, only against the backdrop of the motor lived body (I consider bodiliness as grounded in the motor lived body, in an implicit "I can and do move in this and that way") and of the cognitive-emotional lived body (I consider grabbiness as grounded in the cognitive-emotional lived body, in the subjective experience of the evaluation of the aroused subjective cognitive-emotional lived body) the subject interacting with the environment or with a situation has access to in the process of sensorimotor coupling.

I would like to make the point here that with these changed premises the feel of the emotion "fear" emerges from the sensorimotor coupling of the subject with the environment determined by bodiliness and grabbiness grounded in the motor and cognitive-emotional lived body, which as a backdrop of the actions of the subject's living body in the environment co-determines the subject's sense- making of the situation she interacts with. In this case movement is not instrumental, but constitutive of the feel of an emotion, for every cognitive-emotional evaluation of a situation and the felt sense of an emotion are subordinated to movement in the environment and hence to motion. It is important to remark here that they are subordinated to movement because every emotional evaluation of a situation according to enactivism can take place only if the subject actively explores the environment and actively brings forth – instead of picking up – an own cognitive domain. Grounding bodiliness in the lived body movement becomes constitutive in the determination of the feel of an emotion, for it becomes an implicit "I can and do move in this and that way", which the subject has access to according to the situation she interacts with. From my point of view, in considering bodiliness only as a way to move the body out of the dangerous situation in interpreting the feel of fear O'Regan [16] reiterates once more the instrumental non-enactive role of action and movement.

Against this background, I consider the sensorimotor approach to the feel of experience and emotion as *semi-enactive* because it focuses only on the bodily abilities to engage in interaction and hence only on the objectivity of experience, on the experiencing agent's bodily markers an external observer can objectivize.

4 Conclusion

The aim of this contribution has been to answer the question of how enactive O'Regan's dynamic sensorimotor account of phenomenal consciousness is by focusing on his sensorimotor approach to raw feel. I have argued that O'Regan's approach is *semi-enactive*, for movement and action are just instruments to extract the raw feel of something from the environment we interact with and for the anti- dualistic embodied mind thesis is not taken into account neglecting in this way the question of the realm of consciousness, this is the part of our cognition we access from a subjective point of view.

In illustrating how O'Regan applies his sensorimotor approach to the feel of emotion I have argued that without integrating or taking into account the phenomenological notion of the lived body the sensorimotor approach to raw feel remains incomplete and partial. Taking into account the subjective lived body leads to considering the relation of co-determination of the inner – the embodied backdrop of human experience – and the outer – the environment and the moving living body – and hence the fact that the feel of experience is someone's feel. This requires the development of a method based on a combination of third-person and first-person approach. More precisely, it requires the measure of the relation between objective qualities of interaction and subjective felt qualities of interaction lived through by the agent or subject experiencing the qualities of interaction subjectively. In my view, the investigation and measure of raw feel can be but relational.

References

1. Chemero, A.: Radical Embodied Cognitive Science. The MIT Press, Cambridge (2009)
2. Colombetti, G.: Appraising Valence. Journal of Consciousness Studies 12(8-10), 103–126 (2005)
3. Colombetti, G., Thompson, E.: The Feeling Body: Towards an Enactive Approach to Emotion. In: Overton, W.F., Müller, U., Newman, J. (eds.) Developmental Perspectives on Embodiment and Consciousness, Erlbaum (2008)
4. Colombetti, G.: Enactive Appraisal. Phenomenology and the Cognitive Sciences 6, 527–546 (2007)
5. Fingerhut, J.: The Body and the Experience of Presence. In: Fingerhut, J., Marienberg, S. (eds.) Feelings of Being Alive, pp. 167–199. de GruyterVerlag, Berlin (2012)
6. Fuchs, T.: The Feeling of Being Alive. Organic Foundations of Self-Awareness. In: Fingerhut, J., Marienberg, S. (eds.) Feelings of Being Alive, pp. 151–165. de GruyterVerlag, Berlin (2012)
7. Heft, H.: Affordances and the Body: An Intentional Analysis of Gibson's Ecological Approach to Visual Perception. Journal for the Theory of Social Behavior 19(1), 1–30 (1995)
8. Hutto, D.: Enactivism: Why be radical? In: Sehen und Handeln, S., Bredekamp, H., Krois, J.M. (eds.). AkademieVerlag, Berlin (2011)
9. Hutto, D.: Mental Representation and Consciousness. In: Banks, W.P. (ed.) Encyclopedia of Consciousness, vol. 2, pp. 19–32. Elsevier, Oxford (2009)
10. Hutto, D.: Knowing what? Radical versus Conservative Enactivism. Phenomenology and the Cognitive Sciences 4, 389–405 (2005)
11. Johnson, M.: The Meaning of the Body: Aesthetics of Human Understanding. The Chicago University Press, Chicago (2007)
12. Kirchhoff, M.D.: Anti-representationalism: Not a Well-founded Theory of Cognition. Res. Cogitans 2, 1–34 (2011)

13. Mossio, M., Taraborelli, D.: Action-dependent perceptual invariants: from ecological to sensorimotor approaches. Consciousness and Cognition 17, 1324–1340 (2008)
14. Noë, A.: Action in Perception. MIT Press, Cambridge (2004)
15. O'Regan, K., Noë, A.: A Sensorimotor Account of Vision and Visual Consciousness. Behavioral and Brain Science 24, 939–1031 (2001)
16. O'Regan, K.: Why Red Doesn't Red Sound Like a Bell. MIT Press (2011)
17. Rudrauf, D., Lutz, A., Cosmelli, D., Lachaux, J.P., Le van Quyen, M.: From autopoiesis to neurophenomenology: Francisco Varela's exploration of the biophysics of being. Biol. Res. 36, 21–59 (2003)
18. Scarinzi, A.: Grounding Aesthetic Preference in the Bodily Conditions of Meaning Constitution: Towards an Enactive Approach. The Nordic Journal of Aesthetics 43, 83–103 (2012)
19. Thompson, E.: Mind in Life. Biology, Phenomenology, and the Sciences of the Mind. Harvard University Press, Cambridge (2007)
20. Thompson, E.: Sensorimotor subjectivity and the enactive approach to experience. Phenomenology and the Cognitive Sciences (2005)
21. Todres, L., Galvin, K.: Embodied interpretation: a novel way of evocatively representing meanings in phenomenological research. Qualitative Research 8(5), 568–583 (2008)
22. Torrance, S.: In Search of the Enactive: Introduction to Special Issue on Enactive Experience. Phenomenology and the Cognitive Science, 357–369 (2005)
23. Varela, F., Thompson, E., Rosch, E.: The Embodied Mind. Cognitive Science and Human Experience. The MIT Press, Cambridge (1991)
24. Varela, F., Shear, J.: First-person Methodologies: What, Why, How? Journal of Consciousness Studies 6(2-3), 1–14 (1999)
25. Weber, A., Varela, F.: Life after Kant: Natural purposes and the autopoietic foundations of biological individuality. Phenomenology and the Cognitive Sciences 1, 97–125 (2002)

Experience and Consciousness: Concepts from the Outside in

Stephen Rainey

St. Mary's University College, Twickenham, London

1 Introduction

The 'feel' of driving a Porsche is unlike that of seeing red (O'Regan, J. Noë, A., 2000). Sensorimotor theory and enactivism hold that looking for mechanisms or something 'inside the head' is a mistake in accounting for this. Consciousness does not 'lie behind' experience and action, but rather that it is *in* experience and action. Studying the actions organisms undertake in environments can provide insight into their consciousness and experience. Taking such actions as the locus of study, moreover, can provide greater insight than can studies of mechanisms that drive such interactions. Studying organism-environment interaction in fact provides insight into mechanisms.

However, it seems at least phenomenologically plausible that actions are caused and controlled by consciousness: that consciousness precedes action in significant ways. This is highly contested in sensorimotor theory. This chapter seeks to bolster an argument for that position, and to do so by utilising varied and distinctive philosophical resources. This is considered useful owing to the interdisciplinary nature of much of the sensorimotor field. As a philosopher in the field, it is useful to orient this position in broader reading, thereby shoring up the relevance of philosophy to the field. Phenomenology, whilst not irrelevant to a discussion of how things 'feel' needn't form the central focus of a discussion here. In fact, this chapter will discuss linguistic and epistemological aspects to the nature of experience (broadly construed) so as to provide input to the field of enquiry opened in sensorimotor and enactivist studies.

Noë (2005) suggests that dealing with perception as an interaction between organism and environment is a way to explain direct perception, i.e. where perception is constituted by an organism-environment interaction, no need arises for talk of representational middle-men. I want to press the point that experience and consciousness require more reflection than is readily apparent from these resources. Contemporary sensorimotor theory makes use of some notions that conflate key ideas, such as a basic relation between genus and species concerning 'experience', and this conflation conceals from critique some epistemic concerns. Of key importance in what follows is the epistemological significance of perception to experience (somehow construed) and the role that inference plays in having experiences.

A lot hinges on the notions of 'experience' and 'consciousness' in this discussion. Yet these are by no means simple ideas. On a broadly Kantian account of human

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_6,

experience, neither the sensory nor the purely conceptual are sufficient in accounting for experience and consciousness. That is, human experience and consciousness cannot be said to be wholly determined by either its empirical inputs nor by its own action alone, not wholly given as if ‘from without,’ nor wholly generated ‘from within’ (McDowell, 1994). It is to be argued, in terms of a pragmatic analysis of experience and knowledge, that concepts ought to be thought of as practical abilities and so that an opposition to the role of the conceptual in explaining experience in fact fudges a fundamental, constructive role that the epistemological plays in experience.

The critical strategy here is complex and it is wide-ranging. These features likely make it challenging to the reader. Nevertheless, in order to press the points at issue, and to provide context for where these points come from, this cannot be avoided.

2 The Senses and Concepts

The apparent dichotomy between the senses and spontaneous conceptual activity is a highly germane topic of study given the sensorimotor and enactivist problematic. The relation between the senses and concepts will speak directly to the priority (or not) of consciousness and to the nature of experience. This will be of particular importance with reference to action. In querying the relations between senses and the conceptual, the role of perception in knowledge and experience is a good place to start in opening these issues.

Something to be avoided in providing a story about this role is the conflation of genus and species. For example, in the case of genuine *vs.* merely seeming-seeings, it can be said that the cases of seeings and seeming-seeings can be combined. They can be combined in the sense that, to a certain level of analysis, a seeming is as good as a seeing. Where I am to report the contents of my visual field, a seeming-sight of red and an actual seeing of red is equal. But if we push the scenario further, and enquire into my experience on a broader basis, there must be a difference between seemings and seeings. The difference must be between a genus of seeings and seeming-seeings *vs* a species of actual seeings. The difference must be that which distinguishes between the two as all seeings are seeming-seeings, but seeings have a different role. How so? If there is no 'how' then seeming so might has the same epistemological status as seeing that things are that way. Epistemic matters are exhausted by the genus conditions.

But an epistemic difference between genus and species does make sense in that a seeing rules out a seeming-seeing. The experience itself of a seeing provides a warrant to differentiate between genus and species. As such, species members have epistemic significance the genus members do not. The feel of seeing red is a genus level description. The experience of seeing red ought to be parsed at a species level. This should be a point taken generally, moreover, not merely one about the difference between illusory experience and veridicality. Seemings have no room for criticism – you seem to see red whether or not red is around, but seeings have scope for criticism: do you *actually* see red? The point is about experience in general, with respect to the role knowledge plays in it. For the seeing, *vs.* the seeming, there needs to be the acceptance of the warrant that the seeming stands to the experience as conclusively beyond the genus and a member of the species.

For cognitive science, the genus may be all that's required. Where the concern is with how cognitive systems do what they do, explanatory strategies can work with the genus alone. This is, however, epistemologically insufficient to deal with the realities of visual experience in another sense. Beyond explanatory strategies knowings and believings and seemings versus seeings are of central import to experience in this other sense. As long as this genus/species conflation is made, a confusion will remain in any talk issuing from ongoing discussion. However, this conflation is a first step only. What remains besides this is the constructive role of the conceptual in framing experience *per se*. This is a complicated exposition of a broad area of thought, and so it must be handled carefully. In order to open the issue, an example follows that is designed to build upon and move beyond the simple epistemological point just made.

2.1 Putting Experience Together

One can easily accept that the feel of driving a Porsche is unlike that of seeing red. The feels of the two scenarios, however, needn't be held to be hermetically sealed from one another. For instance, in the case of having driven a *red* Porsche, seeing red might well *recall* the feel of driving that very Porsche. Merely being conscious of red might be understood to have little cognitive impact, but an *experience of red*, the contention here goes, is something beyond simply being conscious of something phenomenal.

This is worth pursuing more in order to make the point clearly. The point to be made is that experience is something quite structured and internally complex, not reducible to conscious events, but reliant upon consciousness. The point will, in the course of the discussion, be shown to support the idea that in fact, consciousness must be thought of as preceding experience, and that the construction of experience is a complex, conceptually-laden affair. To pre-empt somewhat, the contention goes that consciousness is the space of reasons in which the epistemological distinction between the genus-level and species-level is made. This is intended to provide some input to the sensorimotor problematic by prompting further reflection on the notion of consciousness as *in* experience and action. The starting point here deals with the idea of representation.

Representation is not exhaustively defined as pictorial representation. Pictures are types of representation. Representation is not co-extensive with pictures. If I wish to represent the Eiffel Tower pictorially, I can draw it, paint it, photograph it and so on. Each will be a pictorial representation of the Tower – each will be recognisably be a pictorial representation, let us assume. Each pictorial representation will be different from the other, however. We will recognise this as readily as we recognise the content of the diverse images as the Eiffel Tower.

I can label each of my pictures with the legend, “Eiffel Tower.” I do this for one image with a printed label. For another, I write in cursive script. For a third, I bend wire into the shape of the letters L A T O U R E I F F E L and for a fourth I press indentations into the paper in the Braille pattern,

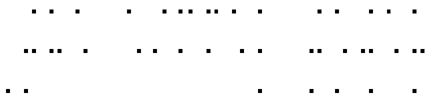

Fig. 1. 'Eiffel Tower' in Braille

Each of the labels is recognisably taking the Eiffel Tower as its referent in some sense. For the last two, knowledge of French and of Braille is required.

While the paintings, photographs and so on are recognisably images of the Tower, the labels are not. However, they are still representations of the Eiffel Tower. Each is unlike the others, but recognisable as representing the Eiffel Tower. of the Eiffel Tower are in this sense *multiply realisable*– one and the same entity can be represented in a range of ways, with diverse materials, for different audiences, with different expectations and consequences. But in each case, the Eiffel Tower is represented. What about the *experience* of the tower itself, besides its representations?

When I stand at the Eiffel Tower, the real, actual Eiffel Tower in Paris, I see before me the brownish hulking structure. It looms into the sky above, it drips water onto my face as I look upward and a chill wind comes from the Seine behind me. I hear a babble of languages and camera clicks and the smell of waffles, crepes and coffee hangs in the air. The visual image of the scene is provided by my looking around, the agglomeration of light stimulating my retina, the memories I form of the things that strike me as salient, the decisions I make to look hither and not thither, and so on.

The light pattern on my retina as I look at the structure before me at time *t* is recognisable as an image of the Eiffel Tower just as any camera image might be. At the point of retinal impact of those particular photons, the analogy could hardly be closer. At this level, the parallel with photography is highly apt, my lookings and my scanning for photo opportunities are tandem processes, one could say.

My olfactory, tactile, auditory experiences likewise are as if printed, cursive, wrought in wire or embossed in Braille representations of the Eiffel Tower. Imagine this: years after Paris, I visit the Statue of Liberty in New York. As I look upward, drizzle dripping on my face, a breeze comes from the waterside. Innumerable languages speak over each other as camera clicks and the smells of fast food and snacks fills the air. “The Eiffel Tower,” I think as for that instant I am brought back to the earlier sightseeing.

The memories I form of the looking around – these are an essential part of my lookings. They are the things I regard as interesting, the things I will tell people about later, the things I would go on to emphasise in any painting I might do later of the Tower. These aren't images as my retinal patterns might be. These are more akin to the labels I make on my pictures. They are not visual, they are not in light. But they are representations of the Eiffel Tower.

The experience of visiting the Statue of Liberty now for me stands as a representation of the experience of the Eiffel Tower experience. *Lady Liberty* needn't be said to be representing the Eiffel Tower, but the sightseeing experience can be said to represent the other. When people ask of my trip I will tell them the two experiences are similar. Clearly they are not in every sense – a brownish, girdered sweep of metal is quite unlike a large greenish lady with a lamp in appearance – but the experience overall of the second sight evokes the first. These experiences are as connected for me the experiencer as "Eiffel Tower," in print, cursive, wire or Braille are to each other and to their respective images in photograph, paint, or anything else. The question remains: how can I relate this to that, one thing as evocative of the next?

An image is a representation, but not one that need be thought of as privileged in some way. When we look 'inside the head,' so to speak, we won't find images, clearly. More generally, if we were to try to find images somehow representing experiences within persons, in their nervous systems, neurons, hearts, livers or anywhere else, we'll be out of luck. The very idea just seems implausible.

When we look in all the places favoured by cognitive scientists (cf. Metzinger, T., 2000 and Jack, A. and Prinz, J., 2003) and to try to correlate brain states and conscious experiences, however, it's reasonable to say that we could find 'representations' everywhere Potential spikes, neuronal activity, magnetic resonance, increased blood flow, facial expressions, movements, gestures, tics, and all manner of things can be put in a category of representations. Just as I can write "Eiffel Tower" in many ways and in many media, I can say there is a brain/neuronal/homeostatic state *s* such that it represents the Eiffel Tower in agent *A*. Just as legends may be written in languages I don't know, or in scripts I can't decode, like French or Braille, it may be a puzzle to read these legends. It would be *unscientific* however to say the lack of a translation manual is the sign of nothing to read. The absence of the Rosetta stone wouldn't mean the hieroglyphs aren't linguistic.

The wet, cold, noisy and smelly reminders of Paris encountered in New York point to the fact that sense modalities aren't readily divisible in lived experience. Just as we might find the representations of experiences distributed across domains we had not expected (neuronal, gestural, bloodflow peaks etc.) so too must it be realised that the representations themselves needn't be univocal. Were we to observe a complex set of peaks and troughs in various parts of the body upon presentation with the Statue of Liberty, we're not in a position to say those peaks and troughs represent the Statue of Liberty. We don't know what they represent yet – could be the smell, the sound, the chill, or indeed it could be the memory of visiting the Eiffel Tower.

What we come to is that the feel of visiting New York is like the feel of visiting Paris. The particular 'give and take' (O'Regan, Noë, 2000) of the venues draw upon the same know-how, despite their description being different in so many respects. Yet there is something in what this give and take provides as an experience that *immanently unproblematically* lends itself to comparing unlike for unlike in similarity. In terms of the epistemological discussion above, there is a genus-species confusion of a sort apparent here: New York seems like Paris, but New York is New York. It is as if one seems to 'see' Paris yet states one sees New York. However, in the more complicated story of experience here over and above perception, the idea of the warrant to say one thing over another is different. The warrant to make claims about the give-and-take of visiting New York draws upon one's experience of give-and-take more generally. The species-level experiences – visiting Paris and New York – draw upon genus-level concepts that knit experience more generally together. Far from being a conclusion in itself, this complicated notion requires further analysis.

3 Concepts: Justification in Action

Whereas it seems straightforward to speak of consciousness and experience, hopefully the foregoing discussions have complicated matters. Moreover, it is hoped that the touristic thought experiment goes some way to demonstrating that, although driving a Porsche and seeing red are felt differently, there is nevertheless a complex story to be told about just what consciousness and experience are in a broader sense. Indeed, we are presented with a predicament as to how to understand the relations between the two. One solution to this predicament (Brandom, 1994) has it that while consciousness might be discussed in terms of non-conceptual elements, experience is a conceptual business.

Noë (2004:199ff) suggests that concepts ought to be considered as abilities and that concepts as such needn't be thought of as overly problematic to sensorimotor theory. Poincaré is quoted as giving an example wherein a ball with differently coloured hemispheres is present. Anyone with the requisite sensorimotor skills, it is argued, will be able to imagine the consequences of that ball's rotation. This is referred to as that individual's possession of the relevant 'observational concepts.'[1] (*ibid*:201) This discussion, however, fails to distinguish between consciousness and experience in at least one way that Brandom elaborates upon.

In Noë (*loc. Cit.*) we don't see how this or that visual representation (say) might prompt the exercise of a particular ability. In the characteristic give and take of any given experience, there are reasons for now giving, now taking. There are reasons to expect this or that. In the case of Poncaré's ball, what is it that might prompt the exercise of the sensorimotor ability of rotating? The point being pressed here is that in experience there are prompts to act this way or that – prompts to act, expectations of give and take. In discussing concepts as abilities, in the abstract, nothing is said of these prompts or warrants issued in experience to set about acting. It is to this that the argument now turns.

3.1 'Nonnatural' Meaning

Let us assume that empiricism is true, at least to the extent that we as human beings require perceptual evidence of contingent matters of fact as a component in

[1] This notion will get specific scrutiny later, with reference to Quine.

belief-formation. Even on this bare account of empiricism, Brandom's position suggests, concepts are at play in the sense that to relate *this* experience evidentially to *that* belief is a matter of concepts. That matter of fact F should prompt belief B is at least a kind of logical inference, and logic draws upon the conceptual in some way. I can be conscious of the conductor's bow without much caring or reacting, but it is a matter of conceptual understanding to experience that bow as *meaning* the end of the concert. The case of meaning is worth dwelling on for a moment, as it will prove informative to draw upon it again later.

As far back as 1957, H. P. Grice gives an analysis of such phrases as 'means', 'means something' and 'means that,' resulting in a twofold categorisation of meaning. In one instance, these phrases can be used in such a way as to indicate *natural* meaning, while in the other they may be used to indicate *nonnatural* meaning (meaningNN). Natural meaning is that sort of meaning involved in statements such as "Smoke means fire", or "Falling pressure means rain". Nonnatural meaning is that sort of meaning involved in statements such as "The flashing lights mean last orders", or "The conductor's bow means the concert is over".

It can be seen from these examples that statements involving natural meaning can readily be paraphrased into a form like "The fact that *p* means that *q*" in such a way that they are incompatible with *q*'s not being the case. It is nonsensical to say "The fact that there is smoke means that there is fire, but there's no fire", or "The fact that barometric pressure is falling means that there will be rain, but there won't be any rain". With cases of nonnatural meaning, this is not the case. It's quite straightforward to say "The flashing lights mean last orders, but it's not last orders", or "The conductor's bow means the concert is over, but the concert isn't over". The lights could be flashing at 8pm due to a dodgy fuse hence not be indicative of the landlord's desire to close, and the conductor could merely have stooped to collect a dropped penny, but there can be no smoke without fire and barometric pressure can't drop without rain resulting.

This demonstrates that statements involving nonnatural meaning cannot be paraphrased into the "The fact that *p* means that *q*" form whilst preserving their meaning. They may be paraphrased into a form that utilises inverted commas, like so "The conductor's bow means 'the concert is over'". The important difference here between natural meaning and nonnatural meaning is that in cases of nonnatural meaning, whatever the part of the statement prior to 'means' is it may be used as an argumentative basis for the part of the statement after 'means'. That's to say, when I state that the conductor's bow means (at least, ought to mean) the concert is over, I'm not saying that her bowing in itself stops the show; her intending to indicate the end by means of a bow does. MeaningNN is thus concerned with intentional, conceptually articulable, deliberate acts on the part of an utterer directed at an audience such that an effect of some kind is expected by the utterer to be evidenced in the audience.

So, Grice gives a three-part intentional account of what it means for a speaker to meanNN something by an utterance for an audience. A speaker's utterance is intended (i1) to elicit an effect in an audience and intends (i2) that the audience recognise her (i1) and intends (i3) that this recognition by the audience of (i1) will feature in the set of reasons that would be given by them for the production of their effect.[2]

Importantly, for Brandom (1994) and in Grice (1957), their conclusions are reached via pragmatic analysis, meaning they are the results of reflection on human beings in action. These methods are akin to those of sensorimotor theory, but deliver on the ideas of experience and consciousness that seem to be implicit in that approach. This pragmatic discussion can feed into contemporary discussion in sensorimotor theory by analysing the very ideas of experience and consciousness upon which so much debate rests. Such a discussion also provides a link to broader philosophical thinking that can provide different perspectives and insights for current thinking.

The case with linguistic meaning, from this sketch of Grice's account of linguistic meaning drawing upon intentions and so on, is clear enough for the purposes of this chapter. This can so far be taken as a point about language. However, the bite is deeper and has epistemological relevance. A linguistic point might be held as reducible to something about conventions (cf. Austin, 1971) or some other phenomenon related to how competent speakers understand one another's utterances. But the really key point we can draw from Grice and relate to Brandom is the role of the conceptual, of the argumentative basis, that appears in the above account.

This points to something about *how we know what we think we know*, and is to this extent an epistemological point. We therefore have to think about the experiential aspects of the conceptual as a route to potential knowledge if we are to unpack the relation between consciousness and experience. This bears upon the conflation of genus and species taken as a starting point in this argument, and begins to look at how that principle operates more generally in terms of the conceptually constructed nature of experience. It is necessary to discuss these points in context before relating them back to the specific sensorimotor problematic under scrutiny here. We must therefore look more closely now at the nature of experience and how experience is understood, over and above the case of linguistic meaning and the understanding of utterances.

[2] Strawson (1971) sees a problem with this characterisation, as it happens, in that a speaker could be thought of as conforming to it while not really being in communication. Suppose I know you to be spying on me, and suppose I also know that you don't know that I know you to be spying on me. Under these circumstances, I could go about arranging evidence that I reckon will lead you to some conclusion such that, were you asked to justify your conclusion at some point, you would mention your covert activities in the justification. In such a scenario, the three conditions of Grice's account are met, but we would hardly call it communication. The difference here is like that between 'trying to get you to think that *p*' and 'trying to let you know that *p*'. The difference may be subtle, but I think it has implications that can be of enough importance to warrant the addition of a fourth condition which Strawson supplies, viz. The speaker must intend (i4) that the audience recognise her intention (i2).

4 Conditions of Experience

A key distinction drawn from the above discussion, and relevant to the one now underway, is that between *reliable differential responsive disposition* and *conceptually articulated response.* In the former case, for example, a reflex might reliably differentiate between states of positive knee-tapping and negative, non-knee-tapping, responding to the positive case through leg-twitching. We may or may not be conscious of this. The experience of having one's reflexes tested at the doctor's surgery, however, calls upon concepts. It does this as the reflex-test case prompts the formation of attitudes in the doctor and patient. The very idea of a reflex response as opposed to a mere leg-twitch illustrates this. The difference is such that various oughts are formed: the doctor ought to carefully aim the reflex hammer so as to make the strike diagnostically significant (not just a whack); the patient ought to generally relax the leg; the leg, *ceteris paribus*, ought to twitch forward upon receiving the hammer; leg-twitching *means* good health, whilst non-twitching indicates something odd is at play. And so on. Throughout all of this we can see the conceptual articulation of the scene as essential to the experience over and above consciousness of phenomena.

We can easily see this toy example as an example of how a particular experience takes on a conceptual dimension, following the trajectory of the earlier Brandom-Grice discussion: the relations of contingent matters of fact to beliefs are mediated in various attitudes and expectations between doctor and patient and between each actor and the world. To labour the point somewhat, the patient expects the doctor will understand the contingent reactions of his body to various stimuli to the extent that she might diagnose problems based on how one thing in the world leads to another, and what belief ought to derive from the complex whole. This permits assessments of health. The doctor expects the patient to be candid, not lying or resisting diagnosis somehow. In these minimal ways at least, doctor and patient are always already utilising conceptual relations of evidence, trust, belief and others in order to have the surgery experience *as an experience* at all. In order that this doesn't seem too much of a leap all at once, it is necessary to turn to Kant (1956) and some of his interpreters in order to clarify the connection being deployed here between the conceptual and experience. Key in what follows is a Kantian idea of concepts as rules.

4.1 Preconditions

Among Kant's contributions to the philosophy of experience can be seen those centring upon the content of the very term 'experience.' W. H. Walsh (1997:42) states it thus:

> […] the 'experience' to which our inferences must be proportioned is experience for a 'mutual us', experience which is shared or shareable, rather than something which is essentially private. But experience of this kind is not so much given as agreed upon; it has to be made out on the basis of what goes on in individual minds. For this process to be carried out we need

> criteria of what is to count as experience proper, criteria which cannot be derived from fact since they determine what it is to be a fact.

The experience to which knowledge must be said to answer is that of public, or shareable experience. In this way, agreement among individuals upon what they experience is construed as the 'touchstone' of knowledge (Kant, 1956:B848).[3] Since these notions of publicity, criticism and discursivity find a place among the criteria for experience proper, experience proper must come from some other source than the antics of nature alone since from these criteria come the bounds of just what can be counted as a fact in our comprehension. Since the objects of the human mind cannot be said to be wholly determined by either its empirical inputs nor by its own action alone, some other factor must come into play in order that it may be said to be determined in some definite respect.

A human mind is one that is passive to the extent that it requires input, but active to the extent that its objects owe something of their constitution to the form of the intellect itself. The objects of our thought are not simply 'given' in experience, so we are not endowed with an intuitive understanding. The objects of our thought are not exclusively intellectual, so we are not endowed with an intellectual intuition. The 'other factor' that comes into play to determine a human mind is judgement. Kant thinks of judgement as the central function of the discursive intellect, hence of the human mind. He writes;

> "[...]we can reduce all acts of the understanding to judgements, and the understanding may therefore be represented as a faculty of judgement." (1956:A69/B94)

Furthermore, Kant goes on to emphasise the role of rules in judgement as follows;

> "Understanding may be regarded as a faculty which secures the unity of appearances by means of rules…" (1956:A302/B359)

The links between understanding, judgement, thought and rules as follows: understanding for a subject is unified through synthesis in judgement according to rules. Rules here are to be thought of as concepts; a concept is here to be characterised as a general kind that serves as a rule for unifying the matter of experience. This is part of Kant's "…general account of concepts of objects, where the concept gives the general form or kind (such as dog or body) that serves as the rule for unifying the matter of sensible intuition; we identify 'substance as phenomenon' (A183=B227) by subsuming the appearances under a count-concept". (O'Shea, 1996:68)

---

[3] Indeed, from the very opening pages of the Critique Kant seems to have a role for discursivity as he lauds his age as "...in especial degree the age of criticism, and to criticism everything must submit."(1956:Axi), while later in the 'Doctrine of Method' he goes on to say; "Reason has no dictatorial authority; its verdict is always simply the agreement of free citizens, of whom each one must be permitted to express, without let or hindrance, his objections or even his veto." (1956:B766 ff)

For Kant, the contents of experience, of experienced reality, must already have a conceptual form, so must already be synthesised according to rules, as the minimal unit of cognitive significance for a discursive intellect is a judgement. An empirical judgement for a Kantian is not a report of a private perceptual state as it might be seen to be for a Humean empiricist. Judgement aims at stating what is the case and, on this characterisation, 'what is the case' is that objectively characterisable order of experience in which many can share. Thus, in order to make claims about an objective order of experience we must make judgements that are objectively satisfied by a common reality.

These abstract conditions are what enable objectively valid, intersubjective discourse to ensue, itself being what enables the possibility of knowledge. The role of what can be called the 'passivity' of the senses and the 'spontaneity' of the intellect is worth dwelling on here, using some more contemporary sources, as it is of key importance in understanding how consciousness might be differentiated from experience as outlined in the beginning of the discussion.

4.2 Concepts, Rules, Reasons

Familiarly within sensorimotor thought, John McDowell (1994) emphasises the role played in Kant by the passivity of the senses, i.e. the sensory faculty, in conjunction with the faculty of the spontaneity of the intellect. Much of his discussions make reference to Kant's famous statement that, "Thoughts without content are empty, intuitions without concepts are blind" (1956:A51/B75) The 'content' here mentioned is empirical content and as such, belongs to the passive sensory faculty. The 'concept' mentioned is the role the faculty of the understanding plays in judgement, and hence belongs to spontaneity.

The main point McDowell wants to establish is that any account of knowledge that overplays or discounts either of these capacities, passivity or spontaneity, is bound to fail. He discusses a natural intellectual 'oscillation' between a coherentist account of truth (one which discounts the passivity of the sensory faculty) and the foundationalist 'Given' (which overplays passivity, discounting the spontaneity of the understanding) against which Wilfred Sellars (1997) famously argues. It is his contention that there is indeed a third way between these alternative schemes, and that this third way echoes Kant's thought.[4]

The idea that Sellars goes against is that somehow there are basic items of knowledge that are simply Given, that is, nonconceptual contents that presuppose no other knowledge and can stand as ultimate groundings for knowledge. McDowell argues that the Given would actually fail for the purpose it is intended anyway. Being non-conceptual, the Given cannot stand as a justification for anything – it is akin to talk at the genus-level, not advancing to the species-level from above. Justifications are matters of concepts.

---

[4] Brandom, more recently, draws upon McDowell and Sellars to make a more subtle argument. Cf. `http://www.nyu.edu/gsas/dept/philo/courses/representation/papers/BrandomNEN.pdf`

One can explain natural happenings exhaustively in physical terms. Justifications, however, require a conceptual articulation since justifications are matters of reasons, hence take an inferential form. The difference is as between the cue ball's striking the black and causing it to enter the corner pocket, and the premise of an argument providing a basis on which to draw a conclusion via a rule of inference. In each case the outcome may be inevitable, but in the former case this is by dint of physics, while in the latter the inevitability is conceptual. The Given, then, can only stand as 'exculpation' for a belief, that is, an admonishment from any possible responsibility for a belief.

McDowell's example of the difference between excuse and exculpation is the imagined case of a person banished from a town.(1994:8n) We are to imagine two scenarios; on one hand, she is found in town illicitly for some reason or other she has, on the other she is in town having been swept up and transported there by a tornado. In the first case, there is transgression of her banishment, hence her illicit presence can be regarded as her responsibility. In the second case, we cannot even justify calling her presence transgression. In the first case she is present where she ought not to be despite her banishment. In the second case she is there because of a tornado.

Rational answerability to experience must consist in reasons and reasons are matters of concepts. Causal accounts of knowledge offer no reasons, only exculpations, particular causes of effects being outside the conceptual arena. To deny this is to go down a road to a causal account of knowledge, or to a conception of knowing minds as being completely determined by the impressions of sense they receive, in which case no question of rationality either internally or externally may arise. Such a denial would amount to an oscillation toward Givenness and, thus, toward exculpations and not justifications or reasons for beliefs. The discerning and use of reasons are of necessity conceptual matters.

The difference between the driving of a Porsche and seeing red, on this account so far, is not that each experience has a particular feel. The difference isn't that each is constituted by a different active knitting of sensory stimulations. The difference is a manifestation of a synthetic judgement such that the possibilities (bounded by expectations) of various sense-modalities provide a nexus of inferential ties between former, present and future states.

The feel of seeing red is unlike the feel of seeing blue in that seeing blue rules out seeing red. More importantly, the experience of seeing red rules out the experience of seeming to see red in being an instance of actual seeing red. However, also as an experience of seeing at all, it is quite redolent with similarities of other seeings and more besides. Seeing Porsche red might easily conjure past experience of Porsche-driving, for instance. The actual experience of such driving, meanwhile, rules out actual Skoda-driving. The give and take of each experience is different – acceleration, handling, confidence – constitutive of what it is like to drive each car is a species-level ability to articulate differences between genus-level drivings. Where my Porsche is damaged, and seems to drive like a Skoda, I am prompted to alter my expectations and, for instance, not to try to overtake as readily. The feel and the experience, though related, are distinct. The former can't constitute the latter, and the latter bears a conceptual structure.

4.3 Some Complications and a Different Approach

Each experienced moment, *qua* judgementally-synthesised and multi-modal framework, provides inferential 'ins' and doxastic 'outs'. These stand as justifications for claims and reasons for expectations. As such, this manifests the difference, in experience, between *reliable differential responsive disposition* and *conceptually articulated response.* So it is too for the difference between consciousness and experience. The consciousness of the experience is in the experience to the extent that the experience is conceptually articulable (as opposed to *mere* consciousness) and so it is derivable from the experience. To recall once more the opening epistemological points, this mirrors the genus-species distinction made with respect to seeming-seeings and seeings. In Kant's case, this derivation takes the form of a transcendental deduction. Importantly for the present discussion, this means that consciousness is appreciably *prior to* experience in being an enabling ground for the possibility of experience. To put it starkly: the conceptual is that something without which the experience couldn't be comprehended *as an experience* at all.[5]

The story doesn't begin and end with Kant and his interpreters, however. Kant uses transcendental reasoning to discover that which must always already underlie experience, or that without which the experience of *x* could not be possible at all. Barry Stroud (1968) rejects the efficacy of the transcendental argument for this purpose as he supposes such reasoning only to tell us what we *must believe* to necessarily underlie what we are investigating. Though we can derive from experience many concepts by various means, they may still lack the 'objective validity' for which Kant strove, meaning that the question of how we gain concepts might be asked and answered entirely with reference to the contingencies of individuals. The modal 'must' of Kant's argument is thereby reduced to a psychological, or subjective necessity of the kind Kant reacted against in Hume (somewhat ironically).

Having begun the present discussion with reference to Brandom and Grice, however, this Stroudian rejection of the strong Kantian ambition for objectivity actually helps rather than hinders the argument. From the outset, it has been the mission of this piece to establish the possibility that consciousness in some respect precedes experience. The strategy has been to show the conceptually articulated nature of experience over and above mere consciousness. In Kant, this was stated starkly through the necessity of judgement in synthesising experience from the twin inputs of the passive senses and the spontaneous intellect.

In departing from the pragmatic stance of Brandom and Grice, however, the use of Kant somewhat talks past the sensorimotor and enactivist point of view. The point being addressed here concerns how consciousness is present in experiences or how it might precede them. It is important now to make the point of the priority of consciousness from a pragmatic point of view. This requires that Kant's position is modified somewhat.

---

[5] With apologies to Jim O'Shea (1996), loosely paraphrased here.

The reason that Kant won't do in the present context is that he relies upon finding ultimate conditions for the possibility of knowledge and experience. In this, he relies upon various notions, among these, the analytic/synthetic distinction. It is to this distinction, and the disentangling of it, that we now turn in order to reassert the pragmatic credentials of the position taken from the outset. Kant gives us the analytic framework and the basis upon which to proceed, but we need a pragmatic stance to make the argument hit home. We need to see what Kant's picture looks like without the analytic/synthetic distinction.

From the beginning, it was argued that a certain genus-level/species-level conflation in the background of sensorimotor theory, in terms of knowledge, meant that an account of experience could not be forthcoming. Through complicated means, arguments have been constructed to try to demonstrate a structure that could accommodate experience. This required a quite hefty account of the conceptual. Exploring an alternative view, the argument now turns to an attempt to account for experience without reference to the genus/species conflation. The beginning point now is a naturalised conception of knowledge which on the face of it ought to be conducive to a sensorimotor account. However, this view will be shown to undermine itself, and in fact lead to a reassertion of the necessity of the conceptual in a manner like that already explored.

4.4 Epistemology Naturalised

The kind of oscillation McDowell considers can be introduced and discussed in terms of some of the aspects of W. V. O. Quine's thought. In 'Two Dogmas of Empiricism' (1980) he famously argues against the empiricist dogmas of analyticity and reductionism. Against the former, it is argued that it is only in degree of willingness to accept or reject statements that purportedly analytic and synthetic sentences differ; none are immune to revision. Against the latter, it is argued that theory faces 'the tribunal of experience' as a whole, there being no foundational experiential reports. The two dogmas are one and the same at root, in effect: verification (the reductionist dogma) uses the idea of sameness of experience to validate sameness of facts. This leans on a difference between fact and language and 'sameness' that itself underwrites the notion of confirming and refuting experiences. Their end-point comes in the idea of analyticity – things that *just are* or not by dint of how their nature.

However, as a committed 'lay physicist' (1980:44) Quine also supposes that the contents of experience are described exhaustively by the 'irreducible postulates' of a physical science:

> "Physical objects are conceptually imported into the situation as convenient intermediaries – not by definition in terms of experience, but simply as irreducible posits comparable, epistemologically, to the gods of Homer... Physical objects and the gods differ only in degree and not in kind."

In Kant, we see the Fact/Language divide addressed (for good or ill) with the notion of synthesis. This is summed up neatly in his phrase, "Thoughts without content are empty, intuitions without concepts are blind. The understanding can intuit nothing, the senses can think nothing." Now, Quine says, "The unit of empirical significance is the whole of science," (1980:43) but owing to the underdetermination of theory by data, in any 'recalcitrant' experience any belief can be revised or kept depending on willingness to revise or keep others.

> "The totality of our so-called knowledge or beliefs ...is a man-made fabric which impinges on experience only along the edges. ... A conflict with experience at the periphery occasions readjustments in the interior of the field. Truth values have to be redistributed over some of our statements. ...But the total field is so undetermined by its boundary conditions, experience, that there is much latitude of choice as to what statements to re-evaluate in the light of any single contrary experience. No particular experiences are linked with any particular statements in the interior of the field, except indirectly through considerations of equilibrium affecting the field as a whole." (Quine, 1980:42)

Crucially: "Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system," and "No statement is immune to revision." Quine also says that recalcitrant experiences *vis à vis* theses objects can stand as reasons to alter theory. Statements concerning experiences of physical objects are supposed to be less theoretical than others in the theory and so are supposed to inhabit the periphery of the theoretical fabric, or field of force, of which they are proposed to, in part, constitute. This account has it both that experiences of objects can stand as reasons to revise theory, but that experiences of physical objects are exhausted by a causal story.

> "A conflict with experience at the periphery occasions readjustments in the interior of the field. Truth values have to be redistributed over some of our statements. Reëvaluation of some statements entails reëvaluation of others, because of their logical interconnections . . . Having reëvaluated one statement we must reëvaluate some others (1980:43)

Quine's arguments against the analytic/synthetic distinction argue against any sense in which theory is constrained 'from within.' Even laws of logic are revisable given, say, experiences of wave/particle duality in physics. However, if theory is answerable in any sense to the facts of physical experience it must be the case that experience is conceptually structured in some fashion, i.e. has rules for unifying the matter of experience.

If really accepted, Quine's position undermines itself in two respects. Firstly, the epistemological model we are asked to accept, that of 'man-made fabric' or 'field of force' at whose centres lie nothing (there are no analytic truths we may not do away with) and whereupon experience impinges only at the periphery, falls down on a truly holistic account, i.e. one which admits no special role for anything analytic. In the face of a recalcitrant experience, we are supposed to re-distribute truth-values across the statements of our theory, all of which are vulnerable to this re-evaluation. No constraint can be placed upon this re-distribution, on pain of inhabiting the necessarily empty centre of the fabric, that is, without analytic rules the idea of an inferential consequence is lost. Jack Arnold and Stewart Shapiro (2007) make just this point;

> "What we still need, I think is an account of inference in the web of belief. I take it that, for Quine, logic guides the process of reevaluating sentences, but what is to guide logic? How should we proceed when logic principles themselves are being reevaluated—when the very extension of "logical consequence" is up for grabs? We need an account of correct inference that somehow transcends the particular logic in place now, and is compatible with a change in logic." (Arnold, J. and Shapiro, S., 2007:11)

Arnold and Shapiro hint at Quine's 'theorem fixation', his conception of logic as consisting in statements and down-playing of the role of rules. They also distinguishes a 'logic friendly' Quine from the more radical early Quine, under discussion here, and go on to say, regarding the passage from 'Two Dogmas' quoted above;

> "How are we to understand the claims that truth values have to be redistributed, that one reevaluation entails others, and that if we reevaluate one statement, we must reevaluate others? The logic-friendly Quine has a ready answer to these questions. The "must" is that of logical consequence. Logic codifies the norms of correct inference that hold in all subject matters, under all circumstances. So logic itself is exempt from the holism. We cannot reevaluate it, since to contradict logic is incoherent. Indeed, this is (part of) what it is to be incoherent. But the radical Quine does not have this answer available. According to him, we are allowed to contemplate a change in logic, if the flow of experience suggests such a move. Again, what does "allowed" mean? That is modal, too. Once again, it cannot mean that it is consistent to change logic, and, say, believe a contradiction, since it clearly isn't." (Arnold, J. and Shapiro, S., 2007:*Loc Cit*)

Things begin to look increasingly arbitrary as the Quinean position is probed. The choices between logics and statements adhered to can't be thought of as having to be made as a consequence of anything unless we are credited with, or bound to, at least some basic analytic operational rules of inference.

Quine (1975) denies a role to the conceptual in experience wholesale. Experience for him just is sensory stimulation by physical objects. This is a self-conscious conflation of the epistemological genus/species distinction made from the outset. This effectively puts seeming-seeings on a par, epistemologically, with seeings. Thus, he has his underdetermination thesis; "If in the face of adverse observations we are free always to choose among various possible adequate… …modifications of our theory, then presumably all possible observations are insufficient to determine theory uniquely." (1975:313)

Without operational rules, analytically valid inferential 'moves', little sense can be made of Quine's metaphor of a field, empty at centre and periphery, and the flow of experience. There is no *transcendental* normative backing for theoretical change either within in terms of operational rules, or without, in terms of answerability to experience. No sense can be made of the idea that any experience can occasion anything if we are to suppose that the very notion of logic itself is up for grabs. Moreover, if we decide somehow that one sentence in our repertoire does need revision, we are still without any reason to suppose that any others might too as a consequence. With logical rules construed as just more statements open to revision, the idea of inferential links between statements becomes empty.

But norms needn't be transcendental to be effective. In fact, from a pragmatic perspective such as that adopted by Brandom and Grice, toward which Quine's thought seems to lead (and the sensorimotor paradigm in general), the norms relevant to comprehending experience and meaning are manifest in action: The conceptual is a realm of rules which constrain possibility through constituting expectations in real people, engaged in real action. These expectations can manifest in simple questions – though questions without simple answers.

Quine's rule-averse position can be seen to tend toward coherentist pole McDowell mentions in the intellectual oscillation he sees. If any adequate theory will do, in the sense that experience is insufficient to determine theory uniquely, the sense in which experience can be said to confront theory at all seems greatly diminished, if not destroyed. The (major) problem with a coherentist account of truth as here sketched is that it can't really accommodate the idea of a rational constraint ‘from without' upon theorising. It discounts passivity just as the account that relies upon Givenness discounts spontaneity.

The coherentist can't easily accommodate the notion of subject matter. They could say; if my only concern is with maintaining the internal coherence of my system of thought, then any changes I may choose to make within that system will be moves in a private game. It isn't necessary that any connection between theory and world need obtain, on such an account. McDowell refers to this predicament as theory's 'frictionless spinning in the void'. In terms of existing sensorimotor theory literature, in this we can see a basis to object with Prinz (2006) and Block (2005) concerning the constituting/causing of experience by agent-environment structure: if constraints 'from without' are missing, the role for that structure is problematised.

The point to be made here is that there is indeed 'friction' acting upon theory and that friction comes from the world. And, moreover, if this friction is to count as a reason to change or maintain elements of the system of thought in question, the content of experience must be thought of as conceptual: it has to provide 'justificatory ins' and 'doxastic outs', to provide give and take. The redness felt, the give and take of the Porsche-driving scene, are the friction, but unless this friction is conceived of as a basis to act one way rather than another, to exercise one ability rather than another, there is no way to account for how these things are experiences rather than fleeting sensory stimulations. Unless these things are thought of in terms of reasons, there is no basis to suggest that an object's apparent increase in size in the visual field constitutes its approaching the agent rather than its expansion. In other words, the know how, the concepts-as-abilities that structure experience over and above mere stimulation, require some kind of inferential form. The diverse sources and wide-ranging nature of the preceding discussions have made for a complex piece of reading.

It will now be informative to synthesise the foregoing into a new form such that the epistemological, experiential, conceptual and other key points might be seen in effect, without so much recourse to textual meanderings. The following therefore stands as a conclusion to the above, as well as a restatement.

5 Knowledge, Action, Being

If we ask, “What does a conscious being do?” We could expect answers like: thinks, feels, acts; Obeys the law; Observes cultural norms; socialises. These are all true and expected. Unexpected but true answers might include... well, what, exactly? Here is a difference in possible descriptions about objects versus conscious beings. An essential part of being conscious is having a constitutive role in the actions one enacts. For instance, someone may try to catch a ball, but fall in a lake instead. Their report that they were trying to catch the ball cannot be ruled out as irrelevant by a third person's statement otherwise.

Maintaining 'falling in the lake' is the relevant action from an outsider point of view doesn't do the same work in the face of the self-report as does such a report of an object. The 'trying to catch the ball' is inherently relevant to the falling in the lake. The experience of ‘a missed catch’ cannot be captured by any but a few physical descriptions of the scene. The intended action is bound up with the possibility of affording the error that is falling in the lake, as it is part of the conceptually articulable features of the scene. A third party's interest in lake-falling doesn't make the catch-effort irrelevant to the context. At the extreme, fixation with lake-falling to the exclusion of catch-efforts misses something essential to this context.

It has thus far been maintained that consciousness precedes experience in some sense, standing as a grounding for it and therefore being derivable from it in some sense. It now seems that self-consciousness appears to be an intrinsically relevant feature of scenarios in which conscious, active beings are present. This is owing to the nature of *pragmatically construed norms*. This point bears some elaboration. In essence, it has been covered: it is isomorphic in certain respects to the preceding discussions. However, in being a subtly new presentation, it needs to be said explicitly.

In perceiving any object as having some quality (x as Q) it is always already implied that the recognition of the possibilities for Q-ing are present in x. We might, for instance, see a ball as 'catchable'. This means that, in somewhat Kantian fashion, perception is equal to a judgement – when we represent x to ourselves as Q-able this is at least in part because we judge that x under suitable descriptions can be Qd.

Furthermore, this means that we employ beliefs, desires and values in order to ground hypothetical modal statements of the form “Possibly If x then Q”. The issuing of such a hypothetical must then have some manner of inferential work, embodied in the perception *qua* judgement, implicit within. How else could Q-ability be understood at all?

Perceiving Q-ability is thus related to belief, desire, value and experience *via* judgement – the judgement that is equal to the perception is not a simple propositional synthesis of subject and predicate but a complex material inferential association — an association of justificatory ins and doxastic outs — between x, ones own knowledge, and possible outcomes based in their conjunction. This suggests that the implicit inferential work within the perception *qua* judgement is *material-inferential*. This materiality of the inferential constellation in play is important as, unlike standard logical inferences, material inference utilises the content of the concept invoked in a predicate, i.e. knowledge of intensional scope – connotation, conversational implicature etc. – pragmatic aspects of competent language use. So the inferential constellation in play is related *via* language to knowledge of logical, physical, technical, practical laws and self-knowledge.

This last factor is important. It is the 'can I φ?' of 'have I got it in me to φ?', where 'φ' is some action. Overcoming character flaws, supererogatory actions, self-effacing behaviour etc. are relevant in assessing one's ability to φ. These factors, indeed, can effectively trump others in practice, i.e. excessive humility may prevent a logically, physically and technically possible action from being carried out.

This highlights an essential characteristic of self-consciousness as an interesting case of consciousness more generally. Agency is a multi-dimensional concept, with rational, linguistic and moral dimensions inter-penetrating one another. What's more, it is essentially self-referential and in a sense self-constituting. This is what underlies the notion from above that suggests self-consciousness is constitutive of a system wherein a conscious being is present. It is worth dwelling a little longer on this notion of 'Can I φ?' and all it presupposes and entails in order to nail down the pragmatic aspects of the discussion, beyond Kant.

Representing x as Q-able is not a matter of picturing, but a matter of associating x with a sphere of reasoning. So perceiving x as Q-able is related to reasons – it has to be as Q-ing x will be an action φ and so must be differentiated from mere movements or reflexes by issuing from a reason base. The reasons in the case of Q-ability are themselves the contents of predicates construed as material inferences in association with the modes of knowledge just expressed (logical, physical, technical, practical and self).

This doesn't have to be a belief that x is Q-able or that 'Yes, I can φ!' as the material-inferential consequences to which I may be committed may yet be unknown to me in their entirety. I may just happen never to have considered that x is Q-able in virtue of its being φ able, for instance (e.g. I consider copper ductile, but happen never to have considered it ductile in virtue of its malleability. In fact, the former is a result of the latter, but it may never have occurred to me) Grasping a concept in material inferential terms needn't be 'complete'.

This doesn't have to be a conscious representation of x as Q-able: the other way around – x can be represented as Q-able consciously in virtue of its being judged as located in a complex of justificatory ins and doxastic outs that are equal to its perception – the perception has been shown to be a judgement, after all. Conscious decisions about x are made *on this basis*, not as the foundation of this basis. The judgement is the perception and it structures possibilities of φ-ing for the agent in advance.

This mirrors the empiricist intuition that objects both precede and constitute perceptual experiences. What it goes beyond, however, is the notion that consciousness is exhausted by *thought*. It does this by identifying the connections between consciousness and experience *qua* nexus of articulable material-inferential states governed by attitudes to pragmatic norms. Moreover, it adds substance to the claim that Porsche-driving feels different to seeing red.

> "The experience of driving a Porsche would seem to consist not so much in the occurrence, as it were all at once, of a certain sensation (something like a twinge or a wave of dizziness), but rather in the characteristic pattern of integrated activity in which the actual driving of the car consists. The particular experience of Porsche-driving comes from the typically Porsche-like give-and-take between you and the car and the things going on around you when you drive the Porsche." (O'Regan, Noë, 2000)

This 'give and take' is no more and no less than a complex of justificatory ins and doxastic outs that structure the experience as an experience at all. That particular complex is what Porsche-driving means, and it's this meaningfulness that makes the difference between driving Porsches and Skodas, driving a hard bargain, seeing red, seeing oneself out and any other bit of human experience. This can be said to be so owing to the adoption of a pragmatic stance and the analysis of human experience as conceptually articulated, cross modal and self-consciously significant if it is to be comprehended as an experience at all. To be cognitively significant experience, indeed, is to be conceptually articulable experience.

References

Arnold, J., Shapiro, S.: Where in the (World Wide) Web of Belief is the Law of Non-contradiction? Noûs 41, 276–297 (2007), doi:10.1111/j.1468-0068.2007.00647.x

Austin, J.L.: Performative-Constative. In: Searle, J.R. (ed.) The Philosophy of Language, pp. 13–22. Oxford University Press, Oxford (1971)

Block, N.: Review of Alva Noë. Action in Perception (2005), http://www.theassc.org/files/assc/Shortened_Noe_Review_JoP.pdf (accessed August 2013)

Brandom, R.: Making it Explicit. Harvard University Press, Cambridge (1994)

Dummett, M.: Frege: Philosophy of Language. Harvard University Press (1981)

Grice, H.P.: Meaning. Philosophical Review 66(3) (1957). Reprinted as ch.14 of Grice. Studies in the Way of Words. Harvard University Press, pp. 213–223 (1989)

Jack, A., Prinz, J.: Searching for a Scientific Experience. Journal of Consciousness Studies 11, 51–56 (2003)

Kant, I.: The Critique of Pure Reason (Kemp-Smith Trans.). St. Martin's Press (1956)

McDowell, J.: Mind and World. Harvard University Press, Cambridge (1994)

Metzinger, T. (ed.): Neural Correlates of Consciousness - Empirical and Conceptual Questions. MIT Press, Cambridge (2000)

Noë, A.: Action in Perception. MIT Press, Cambridge (2005)

O'Regan, J., Noë, A.: What it is like to see: A sensorimotor theory of perceptual experience (2000), http://nivea.psycho.univ-paris5.fr/Synthese/MyinFinal.html (retrieved) (accessed April 30, 2013)

O'Shea, J.: Kantian Matters: The Structure of Permanence. In: Acta Analytica pp. 67–88 (1996)

Prinz, J.: Putting the Brakes on Enactive Perception. Psyche (12.1) (2006), https://www.uoguelph.ca/~abailey/Resources/Prinz,%20Jesse%20-%20Putting%20the%20Brakes%20on%20Enactive%20Perception.pdf (accessed August 2013)

Quine, W.V.O.: 'Two Dogmas of Empiricism' in From a Logical Point of View, pp. 20–47. Harvard University Press (1980)

Quine, W.V.O.: On Empirically Equivalent Systems of the World. Erkenntnis 9, 313–328 (1975)

Sellars, W.: Empiricism and the Philosophy of Mind. Brandom, R. (ed.). Harvard University Press, Cambridge (1997)

Strawson, P.F.: Intention and Convention in Speech Acts. In: Searle, J. (ed.) The Philosophy of Language, p. 27. Oxford University Press, Oxford (1971)

Stroud, B.: Transcendental Arguments. Journal of Philosophy 65(9), 241–256 (1968)

Walsh, W.H.: Kant's Criticism of Metaphysics. Edinburgh University Press (1997)

Sensorimotor Knowledge and the Radical Alternative

Victor Loughlin

Centre for Philosophical Psychology,
University of Antwerp,
Belgium
victorloughlin@hotmail.co.uk

Abstract. Sensorimotor theory claims that what you do and what you know how to do constitutes your visual experience. Central to the theory is the claim that such experience depends on a special kind of knowledge or understanding. I assess this commitment to knowledge in the light of three objections to the theory: the empirical implausibility objection, the learning/post-learning objection and the causal-constitutive objection. I argue that although the theory can respond to the first two objections, its commitment to know-how ultimately renders it vulnerable to the third and arguably most serious objection. I then suggest that sensorimotor theory has two options: concede the causal-constitutive objection or challenge it. I shall argue for the latter. I will claim that a radical sensorimotor theory offers the best means of responding to this objection.

Keywords: O'Regan and Noe, sensorimotor theory, sensorimotor knowledge, know-how, practical understanding, radical enactivism.

1 Introduction

O'Regan and Noe (2001) have argued that what you do and what you know how to do constitutes your visual experience. Sensorimotor theory challenges internalist notions by claiming that it is embodied know-how or skillful engagement with the environment that realizes such experience. This theory has undergone many changes since its inception in 2001 (for example, see Noe and O'Regan, 2002; Noe, 2004, 2009; O'Regan, 2011) yet throughout these changes, proponents of the theory have remained committed to the claim that visual experience is realized by embodied know-how or skillful engagement. Indeed, Noe (2004) states that "[t]his is one of the central claims of the enactive or sensorimotor approach to perception" (p64)

The theory has had some high profile critics. Prinz, Aizawa, Clark and Block have all argued that it faces a number of important objections. In this paper, I will focus on three of these objections: the empirical implausibility objection, the learning/post-learning objection and the causal-constitutive objection. I will argue that although the theory (both in its original 2001 formulation and later incarnations) can respond to the first two objections, its commitment to know-how ultimately renders the theory vulnerable to the third and arguably most serious objection.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_7,

I think this leaves the theory with two options. It could concede the point to the causal-constitutive objection. Sensorimotor theory then becomes a methodological and/or epistemic claim about visual experience. Or the theory could challenge the objection. I shall argue that sensorimotor theorists should endorse option two. I will suggest that they do this by “going radical”. Utilising arguments offered by those who have both criticized and developed sensorimotor theory (for example, Hutto, 2005, and Hutto and Myin, 2013), this paper will describe how a radical version of sensorimotor theory can successfully challenge the causal-constitutive objection.

The layout of this paper is as follows. In section 2, I offer a brief outline of sensorimotor theory. In section 3, I examine three objections that challenge sensorimotor theory and argue that the causal-constitutive objection poses the most serious challenge. In section 4, I sketch out a radical sensorimotor theory.

2 Sensorimotor Theory

O’Regan and Noe (2001) argue that “vision is a mode of exploration of the world that is mediated by knowledge of what we call sensorimotor contingencies” (p940). Sensorimotor contingencies are understood to be relations of lawful dependence between features of an agent’s sensory apparatuses and features of the agent’s environment. These contingencies are sensory since they refer to the agent’s sensory apparatuses (for example, eyes, ears, hands, noses etc) and they are motor since they refer to how those apparatuses react to the environment during movement by the agent.

For example, a sensorimotor contingency unique to human vision is the following:

“If you are looking at the midpoint of a horizontal line, the line will trace out a great arc on the inside of your eyeball. If you now switch your fixation point upwards, the curvature of the line will change; represented on a flattened-out retina, the line would now be curved. In general, straight lines on the retina distort dramatically as the eyes move, somewhat like an image in a distorting mirror” (O’Regan and Noe, 2001, p941).

O’Regan and Noe argue that this demonstrates the lawful dependence between movement of a human sensory apparatus (the eyeball) and a feature of the environment (a horizontal line). They also argue that each human sensory modality - vision, touch, taste, sound and smell - corresponds to a unique set of sensorimotor contingencies (ibid).

The key claim O’Regan and Noe make in their 2001 paper is that it is an agent’s practical knowledge of sensorimotor contingencies - the agent’s sensorimotor know-how - that constitutes the agent’s visual experience.

For example,

“the feeling of seeing a stationary object consists in the knowledge that if you were to move your eye slightly leftwards, the object would shift one way on your retina, but if you were to move your eye rightwards, the object would shift the other way. The knowledge of all such potential movements and their results constitute the perception of stationarity” (O’Regan and Noe, 2001, p949).

It is worth noting that there is nothing inherently controversial in the claim that what an agent does influences what the agent perceives. It is a commonplace to assert that actions help shape and guide perception. Where O'Regan and Noe's sensorimotor theory earns its spurs however is in the constitutive role it assigns to embodied know-how or practical understanding. For sensorimotor theory, knowledge or mastery of sensorimotor contingencies is more than just causally important to visual experience. Rather know-how or skillful engagement with an environment is what constitutes or realizes that experience.

Significantly, Shapiro (2011) identifies two possible interpretations of this claim. On the first weaker interpretation, it is only necessary that an agent have the potential to exercise sensorimotor contingencies. According to this interpretation, "it is important only that one has, sometime in the past, acted on the world in ways that created knowledge of sensorimotor contingencies" (Shapiro, 2011, p168). As Shapiro notes, this interpretation is available when, for example, O'Regan and Noe claim that perception of stationarity is dependent on "the knowledge of all such *potential* movements". On the second stronger interpretation, the agent has to "actually practice those actions that reveal sensorimotor contingencies" (ibid). This interpretation is available when, for example, O'Regan and Noe claim that it is movement or action that reveals the sensorimotor dependence between the eyeball and the horizontal line. These two interpretations of what I will call the knowledge claim will play important and decisive roles in the following section.[1]

3 Objections and Replies

3.1 The Empirical Implausibility Objection

A prominent objection to sensorimotor theory is that it is empirically implausible to think that actions or bodily movements are needed in order to have sensory experience. Prinz (2006) offers examples that suggest this claim is empirically implausible when applied to visual experience and Aizawa (2007) offers an example that suggests it is implausible when applied to tactile and/or auditory experience.

Prinz states that

"[p]erception is not impaired by spinal cord injuries that cause paralysis, by paralysis of eye muscles or brain structures that control them, by atrophy of motor cortex in Lou Gehrig's disease, by destruction of action-control centers in parietal cortex, or in frontal cortex (which are presumably destroyed in many cases of Broca's aphasia)… [I]t is certainly noteworthy that no motor deficits seem to undermine the ability to perceive. There are clear dissociations between perception and action.

[1] One of the ways in which sensorimotor theory has been developed since its inception in 2001 is by applying the theory to sensory modalities other than visual experience. For example, Cooke and Myin (2011) offer a discussion of how a broadly sensorimotor approach is applicable to smell. Applying sensorimotor theory to modalities other than visual perception raises its own set of issues. Nonetheless, I would argue that any applications of the theory that retain a commitment to embodied know-how or skillful engagement will have to address the issues outlined in section 3.

People with motor deficits can see the world, and people with perceptual deficits can act in it" (Prinz, 2006, p10).

Prinz's claim is that people suffering from paralysis of the body still retain the ability to perceive the world around them. In which case, it is empirically implausible to argue, as he suggests sensorimotor theorists do, that visual experience is always dependent on bodily movement. As he puts it, there are important dissociations between perception and action.

Aizawa (2007) recounts an example of someone who experienced awareness of touch and sound during surgery despite the administration of anesthesia and neuromuscular blockades. He describes a 74-year old woman who "recalled that during her operation "1) she felt pain during the incision of the abdomen, 2) she heard the operator say, "It is difficult to remove all tumors because the adhesion is very strong" and 3) she remembered someone had been walking around her"" (Aizawa, 2007, p23). This would seem to demonstrate the empirical implausibility of the claim that an agent must move or act in order to have tactile and/or auditory experience of the world around them.

The empirical implausibility objection has had some high profile advocates and although the original 2001 version of sensorimotor theory primarily focused on visual experience, it is relatively straightforward to see how the Aizawa example could be used to block the application of the theory to other forms of sensory experience. Yet I am going to suggest that O'Regan and Noe's original claim has the means to respond to this objection.

As we have seen, O'Regan and Noe claim that visual experience is constituted by know-how of sensorimotor contingencies. Following Shapiro, we identified two possible interpretations of this claim. On the weak interpretation, visual experience only requires the potential to exercise these sensorimotor contingencies. That is, as long as an agent has exercised the relevant contingencies at some point in the past, then they can obtain the relevant experience. However, on the strong interpretation, the agent must realize the relevant contingencies now through actions or bodily movements in order to have the experience.

I would argue that it is the strong interpretation that the Prinz and Aizawa examples target. For if paralysed individuals can still have experience (visual in the Prinz example, tactile and auditory in the Aizawa example), then this suggests that any strong interpretation of the O'Regan and Noe claim is indeed empirically implausible. An agent need not always act or move in order to have sensory experience.

However, if O'Regan and Noe were to reject the strong interpretation and instead adopt the weak interpretation, then they would have the means to respond to the empirical implausibility objection. For if it is the potential role played by embodied know-how that constitutes sensory experience, then this practical understanding can remain even if an individual is currently unable to move. On this interpretation, it an agent's acquired practical knowledge or understanding of how bodily movement and sensory stimulation depend upon each other that constitutes their experience and not simply their current bodily movement. Hence, a paralysed individual can retain visual experience (and/or tactile and auditory experience) since they possess this acquired know-how. A weak interpretation of the knowledge claim then is not vulnerable to the empirical implausibility objection since such an interpretation does not entail that the agent must currently act or move in order to have sensory experience.

3.2 The Learning/Post-learning Objection

One of the ways in which O'Regan and Noe's sensorimotor theory has been developed is by using it to explain sensory substitution devices. For example, it has been used to explain how agents can gain experience of the world around them via Bach-y-Rita's Tactile Vision Substitution System or TVSS (1972).[2]

TVSS consists of a head or eyeglass mounted camera whose visual output is transduced to trigger an array of vibrators which are placed somewhere on the body of a blind (or blindfolded) subject. After training with the device, during which time the subject moves with the device and learns how movement alters the sensory tactile input, subjects begin to report experiencing objects arrayed in three-dimensional space around them. It is also reported that they are able to make judgments about the number, relative size and position of objects in their environment (Noe, 2004, p26).

Noe (2004) has argued that a TVSS device enables the user to replicate (albeit in a limited way) the sensorimotor interaction that a normal-sighted person would have with their environment. He claims

"[t]actile vision is vision-like because (or to the extent that) there is, as it were, an isomorphism at the sensorimotor level between tactile vision and normal vision. In tactile vision, movements with respect to the environment produce changes in stimulation that are similar in pattern to those encountered during normal vision. The same reservoir of sensorimotor skill is drawn on in both instances" (2004, p27).

In other words, the TVSS user is able to gain vision-like experience because they acquire the embodied know-how of the sensorimotor contingencies that, in a sighted person, normally governs visual interaction with an environment.

However, Clark (2009) challenges this explanation of TVSS. He argues that TVSS is not evidence that the vision-like experience of the TVSS user is realized by sensorimotor contingencies. Clark argues that it is problematic to take "evidence for the role of whole sensorimotor loops in *training and tuning* the neural systems that support conscious perception for evidence of the ongoing role of such loops" (p970, emphasis in original). This is because "nothing in the evidence makes this the case. Perhaps embodied activity is just a causal precondition of setting or re-setting parameters in neural structures that once set and activated, suffice for the experience in question?" (ibid). In other words, there may be an important learning/post-learning distinction (what Clark calls "training and tuning") and it is only during learning to use the TVSS device that sensorimotor contingencies play a crucial or pivotal role.

In response to Clark, I will argue that, as with the empirical implausibility objection, the sensorimotor theorist can reply to a learning/post-learning objection (though such a reply may not save the theorist, as we will soon see).

Given Clark's learning/post-learning distinction, it would seem to follow that in any post-learning phase only internal factors can constitute the vision-like experience

---

[2] For more recent versions of such devices, see O'Regan, 2011.

of the TVSS user. As we have seen, TVSS is a touch-based apparatus (since it consists of an array of vibrators placed somewhere on the body) and consequently principally activates (among other things) the somatosensory cortex in the brain of the user (Noe, 2004, p27). The only internal factor that can be appealed to then in a post-learning phase is the somatosensory cortex. The sensorimotor theorist can thus argue that the internalist needs to explain how and why such cortex can realize vision-like experience. That is, how and why does such cortex support vision-like experience as well as tactile experience?

Hurley and Noe (2003, p145) argue that cortex can acquire visual properties when it is embedded within the particular sensorimotor dynamics characteristic of that modality. If so, then the somatosensory cortex of the TVSS-user can be part of the physical processes that realize vision-like experience because such cortex now defers to the skillful patterns of sensorimotor contingency characteristic of visual experience. The advantage of this explanation for the sensorimotor theorist is two-fold. First, it clarifies how cortex that is associated with touch can, when embedded within the right extended sensorimotor dynamics, also become associated with vision and so explains the experience of the TVSS user. Second, it suggests that even if there is a learning/post-learning phase in the TVSS user's experience, internal factors alone cannot explain this since the internal factor i.e. the activation of somtosensory cortex, remains relatively constant.

Clark however is not swayed by these considerations (see 2009, pp971-972). Moreover, Clark could argue that the weak interpretation of the knowledge claim is in fact compatible with the learning/post-learning distinction. For the weak interpretation only requires that sensorimotor contingencies have been exercised at some point. If the exercise of sensorimotor contingencies were to occur during the learning phase (after which sensorimotor knowledge assumes a potential role), then this is compatible with Clark's claim that during a post-learning phase internal (neural/Central Nervous System) processes assume the pivotal role in TVSS. Clark would need to show how somatosensory cortex can realize visual experience. But if, as he argues, the sensorimotor explanation is only an explanation about the *content* of the TVSS user's experience, then such an explanation does nothing to exclude "standard internalist views about the local (neural) *vehicles* of content" (2009, p971, emphasis added).[3] Thus, a weak interpretation actually supports rather than challenges Clark's learning/post-learning distinction. This point will be developed further in section 3.3. For now it will suffice to note that issues of know-how could potentially be problematic for sensorimotor theory.

3.3 The Causal-Constitutive Objection

Block (2005) claims that "even if perceptual experience depends causally or counterfactually on movement or another form of activity, it does not follow that perceptual experience constitutively involves movement" (p6). He argues that how experience is produced merely reveals the causal basis of perceptual experience and does not reveal what constitutes that experience. This is because "[t]o suppose that the issue is one of how experience can be produced is to shift the topic from a constitutive

[3] Though see the hard problem of content (section 4) for a possible reply to this move by Clark.

issue to a causal issue. Certainly the causal sources of our experience include sensorimotor causal loops, but that does not settle the constitutive question" (ibid).

Recall that for O'Regan and Noe, "vision is a mode of exploration of the world that is mediated by knowledge of what we call sensorimotor contingencies" (2001, p940). The causal-constitutive objection is the claim that how visual experience is produced is a separate issue from what experience is. In other words, determining the causal basis of visual experience does not determine what constitutes that experience. Thus, even if embodied know-how may cause perception, this fails to show that such know-how is part of the metaphysically necessary conditions needed for perception to occur.

In section 2, I introduced Shapiro's distinction between the weak interpretation and the strong interpretation of what I called the knowledge claim. As section 3.1 showed, the strong interpretation is empirically implausible. This leaves the weak interpretation. Yet the weak interpretation looks vulnerable to the causal-constitutive objection.

Consider the following example. O'Regan and Noe claim "the feeling of seeing a stationary object consists in the knowledge that if you were to move your eye slightly leftwards, the object would shift one way on your retina, but if you were to move your eye rightwards, the object would shift the other way" (2001, p949). Yet Block could reply that acquiring the feeling of seeing a stationary object may initially involve engaging sensorimotor know-how i.e. learning how the perception changes with certain bodily movements. But on the weak interpretation, once this visual experience is acquired, then sensorimotor know-how assumes a potential role and a potential role is just that - all it means in reality is that you just see the stationary object. There is in fact no longer any need for embodied know-how.

This opens the door to the causal-constitutive objection. For if sensorimotor know-how is only required during the acquisition of visual experience, then it can only show how such experience is produced and not what constitutes that experience. Moreover, as we saw in section 3.2, the weak interpretation is compatible with the claim that internal (neural/Central Nervous System) processes could be the metaphysically necessary conditions that realize or constitute visual experience. Simply put, sensorimotor know-how may cause perception but it is still the brain and processes within it that constitute it.

The causal-constitutive objection thus amounts to a two-pronged attack on the knowledge claim: (1) sensorimotor know-how only reveals how visual experience is produced and not what constitutes or realizes that experience and (2) an orthodox, internalist view of visual experience is actually compatible with claims about sensorimotor know-how. Since the knowledge claim is central to sensorimotor theory, then this is arguably the most serious objection faced by the theory.

What then are the options for sensorimotor theory? I think there are two possibilities. One is to simply concede the point to the causal-constitutive objection. Since the objection targets the constitutive claims made by sensorimotor theory, then the theory could just drop those commitments. Sensorimotor theory then becomes a methodological claim about how we should investigate visual experience and/or an epistemic claim about how we come to have knowledge about (rather than for) such experience. These claims may still be significant even if they would entail a downgrading of the ontological reach of the theory.

However, a further option is to challenge the causal-constitutive objection. It is this option that I shall argue for. Hutto (2005) and Hutto and Myin (2013) have both criticized and subsequently refined the original claims made by O'Regan and Noe. In the following section, I shall show how Hutto's development of sensorimotor theory and Hutto and Myin's radical enactivism potentially gives sensorimotor theory the necessary firepower to challenge this objection.

4 Going Radical

If sensorimotor theory is to challenge the causal-constitutive objection, then it will need to (1) undermine the picture of visual experience that the objection assumes and (2) show why rival internalist accounts of visual experience – the sort of accounts that would support the causal-constitutive objection - are themselves problematic.

Hutto (2005) argues that sensorimotor theory should abandon the idea that know-how is needed for visual experience. He claims, "the basic *character* of perceptual experience is determined by the features of the different sensory modalities and how they respond to specific objects" (2005, p395, emphasis in original). But he denies that this requires any form of knowledge:

"it is not knowledge – not embodied know-how per se – that gives perceptual experiences their character but *facts about the nature of our embodiment in relation to particular active engagements*. These are facts that we do not know and do not need to know in order to have experiences" (2005, p401, emphasis added).

For example,

"I know that if I take an object from a well lit room to a poorly lit one it will look different. In which ways, I cannot say exactly – even when the object is familiar to me. This does not mean that the way I experience is not dependent upon the appropriate sensorimotor contingencies, only that it is not knowledge of them, at any level, that matters to my perceiving" (2005, p398).

On Hutto's account of sensorimotor theory, visual experience is constituted by sensorimotor contingencies but it is not dependent on knowledge or know-how of those contingencies. I think this abandonment of know-how helps undermine the picture of visual experience that the causal-constitutive objection assumes. For on the Hutto view, how experience is produced is determined by facts about our embodiment in relation to particular active engagements. This view entails that these facts and the determining role they play constitute visual experience. In other words, the "how" and the "what" of visual experience have the same explanans, namely facts about our embodiment.[4] As such, how visual experience is produced is not separable from (and so does tell us something about) what constitutes visual experience. This is a very different picture of experience from that assumed by the causal-constitutive objection.

[4] Rowlands (2010, p78) argues that enacted mind (what I call sensorimotor theory) is in fact a claim about embodiment. I think Hutto's version of sensorimotor theory would agree with this. On the Hutto view, the theory is a claim about the embodied nature of visual experience.

However, this is only the first stage in a possible challenge to the causal-constitutive objection. The second stage is to show why rival internalist views about visual experience are problematic.

Internalism requires that the brain play a privileged role in perception. There are a number of considerations that could support this view. One is that future empirical work will reveal that there are special mechanisms within the brain that ensures that the brain plays this special role. If there are any such mechanisms, then the brain is indeed privileged over any potential bodily interaction with an environment. This possibility, were it confirmed, would challenge any externalism about visual experience, though as Hurley and Noe (2003) have shown, claims about brain mechanisms are compatible with sensorimotor explanations. Nonetheless, since it is arguable that evidence of such a possibility is (at present) not available, I will set-aside this consideration for the moment.

Another consideration available to the internalist is to argue that the brain is privileged over bodily interaction because it processes information about the external world, that is, it "trades in" or "traffics in" informational content about the world. Claims about content are at the heart of the orthodox input-output view of visual processing. Briefly, the story is that information arrives in the brain via the sensory organs and it is then processed according to certain rules or algorithms. This results in some representational state with informational content about the world. This state then leads to further processing and/or signals being sent to the peripheries and possible bodily movement. Such a picture obviously privileges the role of the brain (since this is where the real action happens) and consequently provides support for an internalist view about perception.

However, Hutto and Myin (2013) argue that such an internalist story must to face up to what they call the Hard Problem of Content. This is the problem that "positing informational content is incompatible with explanatory naturalism. The root trouble is that Covariance doesn't Constitute Content" (2013, pxv).

Information-as-covariance is the information revealed when there is a reliable covariance between states of affairs. For example, the rings of a tree reliably co-vary with the age of the tree such that the rings can be used to obtain information about the age of the tree. However, "[a]nything that deserves to be called content has special properties - e.g. truth, reference, implication - that make it logically distinct from and irreducible to mere covariance relations holding between states of affairs. Though covariance is surely scientifically respectable, it isn't able to do the required work of explaining content" (Hutto and Myin, 2013, p67). This entails that states of affairs "do not 'say' or 'mean' anything just in virtue of instantiating covariance relations" (ibid). In other words, information-as-covariance does not constitute information-as-content.

This leads Hutto and Myin to make the following claim:

"[i]f covariance is the only scientifically respectable notion of information that can do the work required by explanatory naturalists, it follows that informational content doesn't exist in nature – or at least it doesn't exist independently from and prior to the existence of certain social practices. If informational content doesn't exist in nature, then cognitive systems don't literally traffic in informational content…[T]here is no naturally occurring informational content in the world" (Hutto and Myin, 2013, pxv).

If Hutto and Myin are right, then the sort of internalist story of visual experience sketched earlier must confront the hard problem. For if the brain is privileged during perception because it "trades in" or "traffics in" informational content about the external world via representational states, then it needs to be shown how this claim can be given a naturalistic explanation. That is, it needs to be shown how the sorts of informational covariances that can be confirmed between brain states and events in the external world can substantiate claims about informational content.

The internalist has a number of options here. One is to argue, "contentful properties exist even if they don't reduce to, or cannot be wholly explained in terms of, covariance relations" (Hutto and Myin, 2013, p68). Such properties might be explained by, for example, "some future physics". Another option is to claim that such properties might be "explanatory primitives – metaphysical extras that might be externally related to covariance properties" (ibid). This "might require us to expand our understanding of the scope of the natural" (Hutto and Myin, 2013, pp68-69).

A further move is to simply deny that covariance doesn't constitute content and show that contentful properties do reduce to covariance properties (Hutto and Myin, 2013, p69). Yet, as Hutto and Myin note, "the metaphysical costs [of this move] will be too heavy for most" (ibid). Alternatively, the internalist could aim to "show that the required notion of information is meatier than covariance but is nonetheless equally naturalistically respectable" (ibid). However, the obvious candidate here – Dretske's indication relations (1988) - seems to go beyond information-as-covariance (Hutto and Myin, 2013, p70).

I will not adjudicate on these various options. I list them merely to show what is involved in confronting the hard problem and the sorts of questions an internalist must answer if they wish to claim, for example, that the brain plays a privileged role because it utilizes representational states with informational content.

I think these various considerations demonstrate how "going radical" can offer sensorimotor theory a way to challenge the causal-constitutive objection. First, the theory abandons the knowledge claim, that is, abandons the claim that embodied know-how is needed for visual experience.[5] This succeeds in undermining the picture of experience assumed by the objection. Second, the theory argues that any internalist view of perception that is committed to informational content must confront the hard problem. A radical sensorimotor theory avoids the hard problem by denying that there are naturally occurring informational contents in the brain for basic perceptual states and arguing instead that although the brain is necessary for such states, it is not privileged over body-world interaction. Thus, contrary to the causal-constitutive objection, a radical sensorimotor theory can retain its ontological commitments since it can affirm the constitutive role assigned to sensorimotor contingencies.

---

[5] If a radical sensorimotor theory abandons the knowledge claim, then how does it explain the visual experience of, say, someone who is paralysed? On the Hutto reading, an agent's visual experience is constituted by facts about the agent's embodiment. Someone who is paralysed remains an embodied agent. The extent to which their embodiment differs from an able bodied person will be the extent to which their visual experience differs from an able bodied person. But, conversely, the extent to which their embodiment remains the same will also be the extent to which their visual experience remains the same.

5 Conclusion

I have examined three objections to sensorimotor theory: the empirical implausibility objection, the learning/post-learning objection and the causal-constitutive objection. I have argued that although the theory can respond to the first two objections, it remains vulnerable to third and most serious objection. I then suggested that this leaves the theory with two options: concede the causal-constitutive objection or challenge it. I argued that sensorimotor theory should endorse the second option. I claimed that it could do this by “going radical”. This involves abandoning the knowledge claim and rejecting informational content as needed for basic perceptual states. Setting aside concerns to do with special mechanisms, I conclude that a radical sensorimotor theory can successfully challenge this objection.

Acknowledgements. Thanks for Erik Myin, Jan Degenaar, David Silverman and Zuzanna Rucinska for their helpful and insightful comments on earlier drafts of this paper. Thanks also to audiences at the ASIB, Sensorimotor Theory Workshop at Goldsmiths and the Radical Enactivism workshop at Ruhr-Universitat Bochum who heard earlier drafts of this paper. This research was made possible by grants from the Research Council of the University of Antwerp and FWO-project GOB5312N.

References

Aizawa, K.: Understanding the embodiment of perception. Journal of Philosophy CIV(1), 5–25 (2007)

Block, N.: Review of Alva Noe, Action in Perception. Journal of Philosophy 102, 259–272 (2005)

Bach-y-Rita, P.: Brain Mechanisms in Sensory Substitution. Academic Press, New York (1972)

Clark, A.: Spreading the joy: Why the machinery of consciousness is still probably inside your head. Mind 118(472), 963–993 (2009)

Cooke, E., Myin, E.: Is Trilled Smell Possible? How the Structure of Olfaction Determines the Phenomenology of Smell. Journal of Consciousness Studies 18(11-12), 59–95 (2011)

Dretske, F.: Explaining Behaviour: Reasons in a World of Causes. MIT Press (1988)

Hurley, S., Noe, A.: Neural Plasticity and Consciousness. Biology and Philosophy 18, 131–168 (2003)

Hutto, D.: Knowing what? Radical versus Conservative Enactivism. Phenomenology and Cognitive Sciences 4(4), 389–405 (2005)

Hutto, D., Myin, E.: Radicalizing Enactivism: Basic Minds without Content. MIT Press, Cambridge (2013)

Noe, A.: Action in Perception. MIT Press (2004)

Noe, A.: Out of Our Heads. Hill and Wang publishing (2009)

Noe, A., O’Regan, K.: On the brain-basis of visual consciousness: a sensorimotor account. In: Noe, A., Thompson, E. (eds.) Vision and Mind, pp. 567–598. The MIT Press, Cambridge (2002)

O'Regan, K., Noe, A.: A sensorimotor account of vision and visual consciousness. Behavioural and Brain Sciences 24(5), 939–1031 (2001)
Prinz, J.: Putting the Brakes on Enactive Perception. Psyche 12(1), 1–19 (2006)
Rowlands, M.: New Science of the Mind. Cambridge. MIT Press (2010)
Shapiro, L.: Embodied Cognition. Routledge, Taylor and Francis Group (2011)

The Problem of Invisible Content

Jack Wadham

Dept. Philosophy, University of Sheffield, Sheffield, UK
j.wadham@sheffield.ac.uk

Abstract. This paper identifies, and attempts to resolve, a serious inconsistency in Alva Noë's theory of perception. I argue that a key feature of Noë's enactivist theory of perception, his claim that perceptual content is 'virtual all the way in', is incompatible with his 'p-properties' account of perspectival content. I will argue that the virtual content thesis implies that p-properties, as characterised by Noë, must be invisible. P-properties play an important role in Noë's theory of perception, and they could not play this role if they were invisible. This problem, the 'problem of invisible contents', must be solved by amending either the virtual content thesis, or Noë's account of perspectival content. At the end of the paper I will argue that the virtual content claim should not be rejected out of hand, and then try to solve the problem of invisible contents by amending Noë's theory of perspectival content.

Keywords: Perception, Alva Noë, grand illusion, change-blindness, inattentional blindness, virtual content, p-properties, invisible contents, appearance patterns.

1 Introduction

Before going into the details of the two claims I take to be incompatible, it is first worth describing briefly Noë's motivations for making the two claims in the first place. I will unpack the two claims in more detail in the next section, but for now I just want to give a rough idea of what they are and why they are important to Noë's theory of vision. Let's focus first on the virtual content claim. Noë introduces the idea of virtual content in response to arguments which are supposed to establish the conclusion that the experience we have of our visual field is in fact a grand visual illusion. There are two arguments that are often given for the grand illusion claim. The first is based on the nature of the retina. Only a relatively small area of the retina, the central foveal region, is able to register visual information in full colour and in sharp detail. The majority of the retina, the peripheral, parafoveal region, can only register low-resolution information, and is unable to register information in full colour. But our perceptual field does not reflect this fact; we see a perceptual field in full colour and sharp detail all the way out to the periphery. The second argument relies on change blindness and inattentional blindness studies. These studies show that people often fail to notice dramatic events in, and changes to, perceptual scenes when

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_8,

their attention is engaged in some task, or disrupted by some experimental trick.[1] These studies are supposed to suggest that we do not see as much of our visual field as we think we do.

Noë argues that the sense we have of our visual scene is not illusory. The 'grand illusion' hypothesis is based on an inaccurate view of our phenomenological commitments:

> It is no part of ordinary phenomenology that we experience the whole [visual field], every bit of it, in consciousness, all at once […] it is no part of our phenomenological commitments that we take ourselves to have all that detail at hand *in a single fixation'* [4, pp.56-7]

The idea is roughly this: if, when we experienced a visual field, we took ourselves to be experiencing all the detail in that field in a single fixation, at a single moment in time, then the phenomenology of ordinary visual experience would be illusory. We couldn't possibly have such access to visual detail as, for example, the aforementioned facts about our retinas demonstrate. But we do not take ourselves to experience the visual field in this way, so our experience is not illusory. Instead, Noë argues, we have *virtual* access to the perceptual field:

> To experience detail virtually, you don't need to have all the detail in your head. All you need is quick and easy access to the relevant detail when you need it […] there's no need to build up a detailed internal model of the world. The world is right there and can serve as "its own best model" [4, p.50]
>
> The visual field – my visual world – is not the field available to the fixed gaze. The visual field, rather, is made available by *looking around*. We look here, then there, and in this way we gain access to the world [4, p.57]

The idea is that we do not take ourselves to have access to the whole of our visual field all at once. Rather, we take ourselves to be capable of accessing the detail around us as and when we need it, using our perceptual skills to examine detail as and when we need to. We take ourselves to be presented with a scene whose detail can be accessed as and when it is required. And we are right to do so; we *can* access the scene's detail as and when we need to. Therefore, Noë argues, there is no grand visual illusion. Our phenomenological commitments accurately reflect the reality of our relation to the visual field. It is not my intention to evaluate this argument; I only mean to describe Noë's motivations for committing to the idea that our visual field is virtually present.

We can now move to Noë's motivation for claiming that p-properties are visible. P-properties are the visual properties an object has from a particular perspective. For example, if I were to look at a circular plate from an oblique angle, the plate would have an elliptical perspectival shape. Noë uses p-properties of this sort to explain our perception of perspective-independent properties, like the circularity of the plate:

---

[1] See Simon and Chabris [1] for a seminal study of inattentional blindness. See Grimes [2] as well as O'Regan, Rensink and Clark [3] for some dramatic illustrations of change-blindness.

> To see a spatial feature such as the size or shape of an object is to explore the way the look of the object varies as we move; it is to keep track of movement dependent changes in p-properties [4, p.84]

The idea is that as we move around our environment, the p-properties of the objects in our visual field will change. For example, if I have my head immediately above my plate, the plate will have a circular p-shape. But if I now sit back gradually in my chair, the plate's p-shape will become gradually more elliptical (because my perspective on the plate is becoming more oblique as I move). Noë thinks that it is because I see such variation of p-properties and because I understand the way in which this variation is a function of my changed perspective that I am able to see perspective-independent properties of objects. To continue with the plate example, we might say that one particular p-property of a plate would not uniquely specify the plate's perspective-independent shape. A circular plate viewed from a certain angle might cause a certain p-shape but equally, a slightly elliptical plate, viewed from a slightly different angle might produce exactly the same p-shape. But when I move, I see a variation in p-shape. And when I move in a particular way, the particular variation in p-shape I see would only have occurred if the plate in question was circular. This is an example of the general claim that movement-relative variations in p-properties uniquely specify perspective-independent properties, and our ability to see the former explains our ability to see the latter. So, for Noë, our perception of p-properties explains our perception of perspective-independent properties. In the next section, I will go into more details about Noë's virtual content story and his p-properties story. In doing so, I hope to set the ground for highlighting the inconsistency between these two aspects of Noë's theory.

2 Virtual Content and Temporal Extension

In the last section, we saw that Noë uses his virtual content idea to explain our ability to see a visual field in full colour and detail all the way out into the periphery. But, as the following quotations should make clear, he extends the explanatory scope of the claim, using it to explain our ability to see whole objects. The argument can be seen as an iteration of the argument we saw in the last section. The idea is that, just like the whole perceptual field, any object that we look at presents us with more detail than can be taken in at a glance (or, more accurately, in a single fixation):

> The content of a perceptual experience is not given all at once the way the content of a picture is given all at once. [4, p.215]
>
> […] you cannot factor experience into an *occurrent* and a *merely virtual* or *potential* part. Pick any candidate for the occurrent factor. Now consider it. It too has hidden facets or aspects. It is present only in its potential. Qualities are available in experiences as possibilities, as potentialities, but not as completed givens. [4, p.217]

To see what Noë means in the second quotation, imagine looking at a small image of someone's face. It will not be possible to take in all the details of this face in a single instant. When Noë talks about 'hidden facets or aspects', he means details that go beyond what you can apprehend in an instant. The idea is that having an experience of some object is not like having a completed picture of that object in one's mind. We should replace this idea with the view that to see an object is to have access to all the detail the object presents:

> We can see that our sense of the perceptual presence of the cat as a whole now does not require us to be committed to the idea that we represent the whole cat in consciousness at once. What it requires, rather, is that we take ourselves to have *access*, now, to the whole cat. The cat, the tomato, the bottle, the detailed scene, all are present perceptually in the sense that they are accessible to us. They are present to perception as accessible. They are, in this sense, virtually present. [4, p.63]

So, to see objects is to have skill-based access to those objects. This is what Noë means when he says that objects are 'virtually present'. Notice how this iteration of the virtual presence argument renders more and more content virtual. In the first section, we saw Noë's motivation for claiming that the periphery of our visual field consists of virtual content. Now, he is claiming that it is not just the peripheries that are virtually present. Even if an object is in the centre of our visual field and is the focus of our attention, Noë still thinks that such an object cannot be apprehended in a single instant. It is, therefore, only virtually present. This gets us to the radical claim that 'the content of experience […] is virtual *all the way in*', that is: 'present thanks to your possession of the skills needed to acquire the relevant information at will' [4, p.134]. We cannot apprehend objects in their entirety in a single moment, but must explore them over an extended period of time. Noë writes that 'a perceptual experience doesn't analyse or break down into the experience of atomic elements or simple features' [4, p.135] and, later, that 'When you peel away the layers of potential and merely virtual presence, you are not left with pure phenomenal content, that which, as it were, is present to your mind now.' [4, p.216]. From these passages we can conclude that, for Noë, all perceptual content is virtual, and it is virtual in the sense that it is accessible (or available) by virtue of the perceptual skills at the disposal of the perceiver.

It is important to emphasise the link between the Noë's virtual content thesis and the claim that perception requires a temporally extended perceptual process. This point will later be crucial to my argument. I hope that discussion of the following passages will emphasise the role of temporal extension in Noë's argument:

> [Experiences] are *activities*, events themselves; they are temporally extended patterns of skilful engagement [4, pp.31-2].
>
> Experience is not something that happens in us. It is something we do; it is a temporally extended process of skilful probing. The world makes itself available to our reach […] Experience has content only thanks to the established dynamics of interaction between perceiver and world. [4, p.216]

Both of these quotations clearly display Noë's commitment to the idea that perceptual experience itself is temporally extended. But the second also tells us something about how this idea is supposed to fit into Noë's notion of availability. The idea is that content is available to us only because we are able to engage in the temporally extended process of skilful interaction with our environment. To use Noë's cat example, we do not 'represent the whole cat in consciousness at once' when we see a cat (4, p.63). Rather, we have the abilities necessary to explore the various facets of the cat over a period of time. And it is this capacity to engage in 'a temporally extended process of skilful probing' that explains the fact that the whole cat is perceptually present to us. So Noë's notion of availability is inextricably tied up with the idea that experience is a *temporally extended* form of skilful activity. I now want to discuss Noë's account of perspectival content, with the eventual ambition of showing that this account is inconsistent with his virtual content claims.

3 P-Properties

Recall that Noë gives the name 'p-properties' to the visible perspectival properties objects have from certain perspectives. Noë gives a more precise definition of p-properties as follows:

> That a plate has a given p-shape is a fact about the plate's shape, one determined by the plate's relation to the location of the perceiver, and to the ambient light. The p-shape is the shape of the patch needed to occlude the object on a plane perpendicular to the line of sight. The p-size of the tree is, in turn a fact about how the trees look, with respect to size, from the location of the perceiver: it is identical to the size of a patch we can imagine drawn on the occlusion plane. [4, p.85].[2]

So the idea is as follows. Suppose you look at a plate from a given position. The particular spatial relation between yourself and the plate will determine the exact p-property that you will see from that location. Imagine that you stayed in a fixed position relative to the plate and held out an opaque material perpendicular to your line of sight. If this occluder were exactly the right shape, it could totally occlude the plate, without occluding anything else in your field of vision. If you managed to construct such an occluder, you would have captured the p-shape of the plate from your particular position.[3]

Noë is emphatic in claiming that, relative to any given position at any given moment, an object has a *single* and *determinate* p-shape/p-size etc. He writes that 'there *is* a single apparent size of an object – namely, the unique way that an object looks with respect to size from a particular position. This is secured by

[2] Noë doesn't explain what he means by the claim that the perceiver's relation to the 'ambient light' also determines an object's p-shape. This point need not concern us here though.

[3] Notice that the size of this occluder could vary. If you made a small occluder and held it up close to your eye, this would have a similar effect to a larger occluder held closer to the plate (assuming there's no difference caused by binocular disparity). This point shouldn't worry Noë; he can still claim that the plate has a single determinate p-shape, since shape is different from size (a square is a square no matter how big or small it is).

phenomenology.' [4, p.84].[4] What we also see in this quotation is the claim that the visibility of p-properties is, for Noë, something that is 'secured by phenomenology'. Noë is emphatic in asserting that we *see* p-properties and he clearly regards this as a claim that accurately describes our phenomenology: 'P-properties are themselves *object of sight*, that is, things that we see. They are visible.' [4, p.83].

4 The Problem of Invisible Contents

We are now in a position to see why Noë's account of p-properties is incompatible with his virtual content claims. Recall that, for Noë, perceptual content is not given to us all at once. It is only because we engage in the extended activity of perceptual exploration that perceptual content becomes available to us. Presumably, this applies to p-properties in the same way that it would apply to non-perspectival content. And if this is right, then we can't, say, take in the p-shape of a person's face in profile all at once. So what? One might ask. If I hold still, and look at the person's face over an extended period of time, I can see the face's p-shape, since there will be a single p-shape visible to me for as long as I hold still, and I can take it in at my leisure. But suppose I don't stay still (Noë is keen to emphasise the active, movement-involving nature of everyday perception). If I am moving around while looking at the face, the p-shape visible to me from my particular perspective will change at every instant. When I move around the face, the 'here' (the location from which I am viewing it) will be different at every 'now' (the time at which I see the face). And so, at every 'now', a different face-p-shape will be visible to me. But if content is not given to me all at once, I can't see anything that is only visible for the duration of any 'now'. If content is not given to me in an instant, I cannot see some feature of the world that is only visible to me for an instant. And when I am moving, a given p-property will, by necessity, only be visible to me for an instant. From this it follows that, while in motion, we cannot see individual p-properties.

Notice that my argument above assumes that it is only possible to see a given p-property of a given object when occupying the perspective at which the object presents that p-property. The idea was that, while moving, we can only occupy a given determinate perspective on an object for a brief instant. And, given that p-properties are only visible from such determinate perspectives, this means that any given p-property will only be visible for a brief instant (the instant at which we occupy some determinate perspective). But if we can't take in any content in an instant, then we can't see a property of the object that is presented to us only for an instant. This argument assumes that we can only see a given p-property of an object while occupying the precise location at which the object presents that p-property. This assumption seems plausible and, as the following quotations should establish, it is one that is shared by Noë:

---

[4] For a similar 'occlusion' theory of perspectival features in depiction theory, see John Hyman [5].

> Elliptical is just how circular plates viewed from an angle look. Indeed, we experience the plate *as circular* precisely because we encounter its elliptical look from here, and we understand the transformations the elliptical apparent shape (aspect) would undergo as we move. [3, p.78]

The key point I want to draw out of this quotation is that, for Noë, we experience some particular elliptical p-shape that the plate presents only because we occupy a position at which the plate presents that particular p-shape. So Noë shares the assumption upon which my argument against him rests.

We can now see why my argument to the effect that p-properties must be invisible while we are moving poses a serious problem for Noë's theory of perception. If we can't see p-properties while we move, this undermines Noë's explanation of the fact that we can see the perspective-independent properties of objects. Noë's explanation of this fact is that, as we move, we experience a variety of p-properties presented by the object we are looking at. It is through implicitly understanding the way in which the p-properties we see vary as a function of our movement that we come to see the actual, perspective-independent properties of objects. But if we can't see p-properties while moving, this explanation doesn't get off the ground. Noë wants to explain our perception of perspective-independent properties by appealing to our understanding of the way in which perspectival properties vary as a function of our movement. But if he is going to do so, he's going to have to give an account of our perception of perspectival properties such that these properties are actually visible while we are in motion. For ease of reference, I will henceforth call this problem **the problem of invisible contents**.

The problem of invisible contents arises from the conjunction of Noë's virtual content story and his claims about p-properties. So we have a choice. We can either eliminate the problem by revising Noë's p-properties story, or we can alter his virtual content thesis. I will choose to revise Noë's p-properties story. One might think this is a bad choice. It might seem that Noë's virtual content story is wildly implausible, whereas what he says about perspectival content sounds pretty sensible. Before arguing for an alteration of the p-properties story, I will consider, in order to reject, an argument for the conclusion that Noë's virtual content claim is wildly implausible.

5 Should We Drop Virtual Content? An Argument from Subitization

To see how the argument from subitization works, we first need to get clear on what subitization is. Coined by Kaufman et al. [6], subitization is the ability to enumerate grouped objects that are presented only for a fraction of a second. For example, if I showed you a picture of the side of a die with five dots on it for only a fraction of a second, you would be able to tell me the number of dots on the die. You can do this without counting the dots; you take in the information in an instant. Subitization is just one example of our ability to take in perceptual content in an instant. But the fact that we have such an ability seems to run counter to Noë's claim that the 'content of a perceptual experience is not given all at once the way the content of a picture is given all at once' [4, p.215]. If I can see at once the number of spots on a die, then it looks like this is an example of perceptual content that *is* given all at once; it looks like subitization is a counterexample to Noë's claim that content is virtual all the way in.

But we can avoid this worry by making a very minor alteration to Noë's claim that content is virtual all the way in. We can say instead that content is virtual *almost* all the way in. The idea would be that most visible objects and properties present us with more detail than we can apprehend in a single instant. We may be able to take in extremely simple content, like the number of dots on a die. But most content is more complex, and can only be taken in over an extended period of time. Pick any object in your immediate surroundings. How many of its properties can be taken in at a single glance? Does its colour look uniform, or are some patches better lit, or more in the shade than others? Can its entire shape, or its entire p-shape, be apprehended in a single fixation? Take, for example, the p-shape of a person's head in profile. Such a p-shape is far too complex to be taken in instantaneously, yet Noë's theory requires that we are able to take in such p-shapes in an instant. My conjecture is that most p-properties are, like the p-shape of a head in profile, too rich and complicated to be taken in at a glance, in a single fixation.

So much for the subitization worry. But one might think that the subitization results give Noë a potential solution to the problem of invisible contents. Perhaps p-properties, like small groupings of dots, are the kind of thing that can be apprehended instantaneously. If this is so, then Noë's virtual content claim might be consistent with the visibility of p-properties after all. But it is doubtful that any more than a small number of all p-properties are simple enough to be apprehended in an instant. It might be plausible to think that, say, the elliptical p-shape of a coin held at a certain angle may be apprehended in an instant. Similarly, the coin's non-perspectival shape could probably be apprehended instantaneously. But what about the p-shape of a person's head in profile, or that of a chair, viewed from a certain angle? It is doubtful that such p-shapes can be apprehended in an instant. One explanation of the fact that figurative painting is so difficult is that p-properties are often very complex, and therefore difficult to render. In general, p-properties are no less detailed than non-perspectival properties, and there is no reason to think that, in general, detailed p-properties are any more amenable to instantaneous apprehension than non-perspectival properties.

While Noë's virtual content claim might seem implausible, my altered version of the claim is not so easy to refute. The fact that some features can be apprehended instantaneously does not falsify this modified virtual content claim. But this alteration of the virtual content claim will not give Noë a solution to the problem of invisible contents. Even my weaker version of the virtual content claim renders most p-properties invisible at a single glance. And if this is true, then most p-properties will be invisible while we are in motion. And if most p-properties are invisible while we are in motion, then Noë's explanation of the perception of non-perspectival properties will fail in most cases. And note also that even if it were the case that most p-properties can be apprehended in an instant, those few cases which could not be so apprehended would still cause a problem for Noë's theory.

6 Appearance Patterns: Solving the Problem of Invisible Contents

Earlier I argued that what Noë says about p-properties is incompatible with the claim that perceptual content cannot be taken in at an instant. I argued that while we are moving, individual p-properties will only be visible for an instant. And if we accept the virtual content claim (the idea that hardly any content can be apprehended in an instant), it follows that we can hardly take in any p-properties while we are moving. This is the problem of invisible contents. I now want briefly to sketch an alternative to Noë's p-property story: the appearance-pattern theory. This theory, when combined with the virtual content claim, will not give rise to the problem of invisible contents.

Appearance-patterns are just patterns composed of multiple p-properties. They are temporally extended sequences of p-properties. Because they are just sequences of p-properties, appearance patterns are still perspectival properties. But according to my appearance-pattern view, we can see appearance patterns (patterns of p-properties) even though we cannot see the individual p-properties of which these appearance patterns are composed. To see how this idea works, consider a cinematic analogy. A film reel can depict an event that occurs over time by presenting a succession of individual picture frames, each depicting a visual scene as it is at an instant in time. If the frame rate at which the film is played is sufficiently high, it will be impossible for a viewer to see any one of the frames that of which the film is composed. If the film was composed of a thousand frames per second, for example, each frame would only be presented for one thousandth of a second. Humans lack the perceptual acuity to take in something that is only presented for this amount of time. Nonetheless, they would still be able to take in the overall scene depicted by the film. In other words, they would be able to see that which is depicted by the overall pattern made up of many frames presented in quick succession, without being able to see any of the individual frames of which the pattern is composed. In this analogy, the film-reel was the analogue of appearance-patterns, and the individual frames that composed this film reel were the analogue of the p-properties that compose an appearance-pattern. Like a film reel, an appearance-pattern can only be taken in over a period of time. And just as taking in a film does not require that we see the individual frames of which the film is composed, so seeing appearance-patterns does not require that we see the individual p-properties of which an appearance-pattern is composed.

One might object that although one would not be able to see a single frame of the film if it was quickly presented on its own, there is a sense in which one is able to see individual frames by seeing the general frame-pattern of which they form a part. Similarly, although it is impossible to see a single grain of sand at fifty metres distance, it is possible at this distance to see a sand castle, composed of many grains of sand. One might think that the ability to see the sand castle confers on you the ability to see the sand-grains of which it is composed. I find this counter-intuitive, but even if it were plausible, the thought would not cause a problem for my argument. If the ability to see the castle confers the ability to see the sand grains, then the ability to see the castle has explanatory priority. Likewise, in the film case, it would be the ability to see the film that explains the fact that there is sense in which we can see the individual frames. Put in general terms, the ability to see the whole explains the ability to see the parts. In the case of appearance-patterns, it would be the ability to see the pattern of p-properties which explains the (anyway questionable) sense in which it is possible to see the individual p-properties that make up this pattern. This order of explanatory priority contrasts with that in Noë's theory, according to which our ability to see patterns of variation in p-properties is explained in terms of our ability to see individual p-properties.

Let's now consider how appearance patterns might do the work required of p-properties in Noë's theory. Recall that our ability to see p-properties was supposed to explain our ability to see non-perspectival properties. We see p-properties, and we see and understand how p-properties change as a function of movement. I aim to give a similar explanation in terms of appearance-patterns. As I change my position relative to a plate, I will see a pattern of p-shape-change. The plate's perspectival shape will go from, say, round to elliptical. I can see this pattern of p-shape-change without seeing any individual p-shape of which the pattern is composed. And the fact that the particular way in which I changed position gave rise to the particular pattern of p-shapes I saw is explained by the fact that the plate is in fact circular. If I implicitly understand that this is the case, I can get from appearance-patterns to non-perspectival properties. So we can still say that to see a non-perspectival property, we have to see and understand how perspectival properties (appearance-patterns) vary as a function of movement. While this explanation is similar to Noë's, it doesn't rely on the claim that p-properties are visible. For this reason, it does not give rise to the problem of invisible contents. More will have to be said about exactly how this account is going to work. My intention here was only to sketch a possible solution to the problem of invisible contents.

References

1. Simons, D.J., Chabris, C.F.: Gorillas in our midst: Sustained inattentional blindness for dynamic events. Perception 28(9), 1059–1074 (1999)
2. Grimes, J.: On the failure to detect changes in scenes across saccades. In: Akins, K. (ed.) Vancouver Studies in Cognitive Science. Perception, vol. 2, pp. 89–110. Oxford University Press, New York (1996)
3. O'Regan, J.K., Rensink, R.A., Clark, J.J.: Change-blindness as a result of "mudsplashes". Nature 398, 34 (1999)
4. Noë, A.: Action in Perception. MIT Press, Cambridge (2004)
5. Hyman, J.: The Objective Eye: Colour, Form and Reality in the Theory of Art. Chicago University Press, Chicago (2006)
6. Kaufman, E.L., Lord, M.W., Reese, T.W., Volkmann, J.: The discrimination of visual number. American Journal of Psychology 62(4), 498–525 (1949)

Beyond Vision: Extending the Scope of a Sensorimotor Account of Perception

Caroline Lyon

Adaptive Systems Research Group, University of Hertfordshire, Hatfield, UK
C.Lyon@herts.ac.uk

Abstract. We examine the scope of some sensorimotor accounts of perception, and their application in developmental robotics. Current interest in sensorimotor theories, and the enactive paradigm, was stimulated by the seminal book *The Embodied Mind* by Varela, Thompson and Rosch (1991) [32]. However, both in this initial book and subsequently there has been much work on visual perception and less attention to other perceptual modalities. We suggest that the insights gained from an exploration of the visual domain need supplementing, and in some respects qualifying: some significant characteristics of vision do not hold for audition, in particular for the perception of speech. This leads into a discussion of the importance of integrating different perceptual modes, with particular reference to robots and human-robot interaction. We examine the effect of including audition in accounts of perception, and suggest that it makes sense to avoid the unnecessary straight jacket of a model based primarily on vision and touch alone. The sensorimotor approach can be extended to other perceptual modes.

Keywords: sensorimotor, perception, vision, audition, human-robot interaction.

1 Introduction

In this chapter we examine the scope of some sensorimotor accounts of perception, and their application in developmental robotics. The current interest in sensorimotor theories, and the enactive paradigm, was stimulated by the seminal book *The Embodied Mind* by Varela, Thompson and Rosch (1991) [32]. However, both in this initial book and subsequently there has been much work on visual perception (for example [25,24]), some on touch, but little attention to other perceptual modalities. Varela et al. write of perception in general, but focus on vision and explore in detail the perception of color. They propose that

> color provides a paradigm of a cognitive domain that is neither pregiven nor represented but rather experiential and enacted ..... The time has come, however, to step back and consider some of the lessons this cognitive domain provides for our understanding of perception and cognition in general. [32, page 171]

We suggest that the insights gained from an exploration of the visual domain need supplementing, and in some respects qualifying: when we examine auditory perception we find that some significant characteristics of vision do not hold for audition.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_9,

This leads into a discussion of the importance of integrating different perceptual modes, with particular reference to human-robot interaction. Starting from the philosophical origins of sensorimotor theory we pick up some of the ideas that turn out to be relevant to present day issues. We examine the effect of including audio and other perceptual modes in accounts of perception, and suggest that it makes sense to avoid the unnecessary straight jacket of a model based primarily vision and touch. The sensorimotor approach can be extended to other perceptual modes.

2 Background

The genesis of ideas expounded by Varela et al., and their followers, can be traced back through European philosophers Merleau-Ponty, Heidegger, to Husserl [19,9,12]. (For an accesible overview see [18].) They were also influenced by strands of Buddhist thought. The interest in vision was characteristic of all these philosophers, and a similar trend was evident in the empiricist British school, typified by works such as those of Berkeley [4], Locke [16] and, to some extent, Hume [11]. For instance, Berkeley produced his *New Theory of Vision*, where "new" was 1709. Locke wrote of "sight, the most comprehensive of all our senses ...".

Thus Varela, Thompson and Rosch continued in a field which had given a preeminent position to vision. They set out to counter cognitivist approaches that were influential in the latter 20th century - representational theories that typically proposed some inner picture or symbol mediating between the outside world and the mind. Their insights into the enactive, embodied nature of perception entailed a different account of visual perception.

Varela et al. acknowledge their debt to Merleau-Ponty, who developed his ideas through critiques of the phenomenology of Husserl and Heidegger [19,20]. Merleau-Ponty explained perception as an embodied activity through which we relate to things in the world around us. As he says in *Phenomenology of Perception* "Perception opens a window onto things. This means it is directed, quasi-teleologically, towards *truth in itself* in which the reason underlying all appearances is to be found" ([19], his italics). Though he goes on to talk about perception in general terms "a window onto things" imply vision and touch, not including hearing, tasting, smelling.

However, without pursuing the question of what *truth in itself* might mean, we can see how such a philosphy begins to be relevant to the development of artificial cognition in robotics. There is no homunculus or inner man viewing percepts that are reconstituted as a model of some part of the environment: the subject is inseparable from the world and "the world is not what I think but what I live through" (ibid, page xviii). This need not only mean a world of visible and tangible "things": different perceptual modalities can, in theory and sometimes in practice, be implemented in a robot and integrated to simulate human cognition [22,33]. Examples can be found in work done in the ITALK project, *Integration and Transfer of Action and Language Knowledge in Robots*, described in [5], in which elements of language are acquired by a humanoid robot interacting with naive human participants, through its own sensorimotor experiences - visual, auditory and proprioceptive. Another example is work done in the SPARK project, described in *Spatial Temporal Patterns for Action-Oriented*

Perception in Roving Robots [1] in which the “agent transforms sensory signals to give rise to motor output ... there is no need for an internal model. Perception is active” (ibid, page viii).

The Focus on Vision

Merleau-Ponty provides the starting point for Varela’s philosophy and subsequent developments in sensorimotor theories, which have become focused predominantly on visual perception. In “The Embodied Mind” [32] Varela et al. take color as a case study. Their illuminating account reviews many experimental results showing how the perception of color is a perceived attribute, partially dependent on the observer and on ambient conditions. The color of an object is seen as part of “a patchwork of visual modalities” including size, shape, motion lighting conditions, etc. (ibid, page 162). Interpretations and associations of color are deeply rooted in our culture, and Varela explores specific cognitive processes related to it[1]. However, though they write of perception in a general sense Varela et al have almost nothing to say about other modes of perception apart from a passing reference to hearing and a paragraph on olfaction. In other work from the Enactivist school the sense of touch is explored, and Noë goes so far as to say that “Touch, not vision, should be our model for perception” [23]. Here we again have perception of “things” that could be seen as well as touched, but excluding audition.

The focus on visual perception is indicated by the titles of writings. For instance, *A sensorimotor account of vision and visual consciousness* by Regan and Noë [25] has been very influential. In the preface to a collection of readings edited by Noë and Thompson entitled *Vision and Mind* [24] the editors say ”The writings in this volume investigate the nature of visual perception. Our goal has been to produce a collection that can serve as a starting point for the philosophy of perception.” We argue that the study of visual perception is indeed a starting point.

However, other modes of perception are often integral parts of the perceptive process, and some of their characteristics, explored below, differ from those of vision. Visual processing is only part of the story.

3 The Need to Integrate Multiple Modes of Perception

Though vision is a key mode of perception, in humans and other animals it is critically integrated with hearing, touching, tasting, smelling as well as with internal proprioceptive information. Vision can often not be disassociated from other perceptual modes and the need to integrate them has underscored much recent work in robotics. For example, in work on language acquisition through interaction between humans and a humanoid robot an acoustic sound stream, visual percepts and proprioceptive information have to be integrated [30]. See Figure 1. Vision plays a significant role in language acquisition

[1] A striking anecdote related to the author concerns a dictat during the Cultural Revolution in China. At that time red was the color of revolution and progress, so it was deemed incorrect to have red mean “stop” on traffic lights. A decree went out that red should mean “go” and green should mean “stop”. Chaos ensued until even the most committed revolutionary agreed to reverse the order.

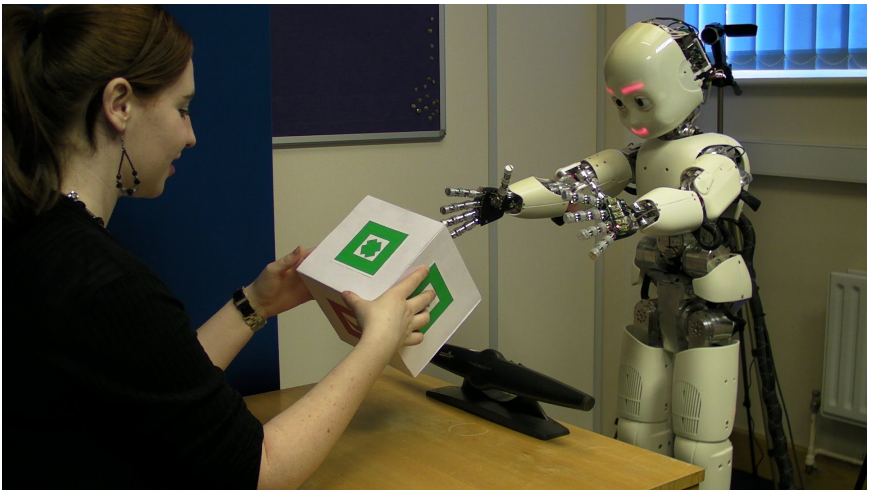

Fig. 1. Experiment with the iCub robot in which a participant teaches it the names of shapes and colors, using visual, auditory and proprioceptive modes of perception

(though not an essential one, as people born blind can learn to speak). For example, research has shown how infants are aware of the gaze of a carer, and shared gaze contributes to language learning.

There are also examples of the integration of different perceptual modes for non-speech sounds. Experiments have been carried out to investigate *vision - action - sound* brain functions using fMRI technology [26]. In observing a familiar action that produces a sound, such as drumming, we can usually predict the sound, but when the natural synchronization was disrupted this had a measurable effect. A second set of experiments investigated the effect on brain processes of disrupting the natural covariance between the velocity of the drumming action and the sound intensity. In both cases expert drummers were compared to non-musicians, and the effects were found to differ. A significant finding was that in synchronised drumming the brain activity in experts was relatively reduced. It was suggested that “the reduced activity found in certain brain areas of musicians was accompanied by an increase of activity in other areas” (ibid, page 1490).

Other examples of the need to study perceptual integration are ubiquitous in robotics. For instance, in modelling a grasping action hand-eye co-ordination has to be mastered [31]. The biological inspiration behind some robotic research has looked to non-human models and in exploring perceptual machines inspired by insects many multimodal interactions between different sensory systems have to be studied: visual, auditory, olfactory, mechanosensory [1].

It is also interesting to consider neuromotor prosthesis studies, such as work on the control of a robotic limb replacing an amputated arm [8]. This is a prime example of a sensorimotor system, based on the capture of neural codes in the subject, and the creation of links between these neural signals and action in the outside world. With closed-loop control the neural activity of the subject guides a device and receives sensory feedback. The perception of this sensory feedback in turn determines the next step in an interactive cycle controlling the movement of the robotic limb. Perception in this case is highly directed, intentional, and active, integrating visual, haptic and proprioceptive percepts.

In a critique of human manual activity Hutto questions whether intentional instructions can control hand movements [13]. “Only very fine-grained instructions would be capable of directing or controlling specific acts of manual activity successfully. This raises a number of questions. How do brains decide which general kind of motor act, M, is the appropriate sort of motor act to use in the situation at hand?”. The control of robotic limbs provides an existential answer to this question: it is through the multisensory feedback cycle.

Tracing back to the ideas of Husserl and Heidigger we can see a link from some of the ideas they examined to issues for neuromotor prosthesis. Grasping actions that are normally done without thinking have to be executed with directed concentration, which is crucial for successful execution. The aim is to move beyond this stage so that actions at a low level become routine, subconscious processes. Heidigger identified these different modes, comparing the routine, subconscious use of a hammer by an experienced carpenter with conscious, directed attention.

Husserl emphasized “intentionality” as a key component of perception. As Dreyfus says, Husserl made intentionality one of the main topics of philosophy [18, page 256]. Now, it is not uncommon in this field to find confusion generated by ambiguous words, and we need to examine the word “intentional”. As well as meaning “directed attention” this also can mean doing something planned [32, page 16], and it is this second meaning which has wider currency in the world at large; for instance, common legal definitions of crimes may include the requirement that they are “intentional” acts. This second meaning is also a crucial concept in neuroscientific modelling (for example in speech production, e.g. [10]) and in robotic development (for example trajectory planning for grasping e.g. [31]) where goal-oriented, forward plans are required. We need to be aware whether “intentional” is used in a present temporal frame as “directed” or in a future temporal frame as “planned”.

Active Visual Processing and Passive Perception in Other Modalities

A key characteristic of visual processing is that it is typically an *active* process. As Varela puts it: ”[t]he enactive approach underscores the importance of two interrelated points: 1) perception consists of perceptually guided action and 2) cognitive structures emerge from the recurrent sensorimotor patterns that enable action to be perceptually guided [32, page 173]. Noë summarises the position succintly as he says: “Perception is not something that happens to us, or in us. It is something we do.” [23]. Perception is a kind of thoughtful activity.

Passive perception can occur occasionally in vision: for instance, if an unexpected flash of light occurs. Passive perception can be exploited in subliminal advertising. However these examples are rare, and typical visual processing is active as articulated by Varela, Noë and others.

Now, contrast this with auditory perception. In this case the hearer can often have a passive as well as an active role. Active listening occurs when a subject directs their attention to speech or other sounds, so humanoid robots acquiring linguistic skills through interaction with a human will typically be listening in an active mode [17]. On the other hand, auditory perception can be passive: a sleeping person can be woken by sounds, and some loud or incessant noise can be an unpleasant experience that the hearer cannot escape. Similarly, perceptions of touch can be active, as in feeling an object, but also passive as in receiving a blow. Perceptions of smell can be active, as in sniffing out the source of an odour, but also passive as when a smell unexpectedly hits your senses.

A striking example of passive perception comes from experiments on the perception of speech sounds as reflected by mismatch negativity (MMN) signals, a component of auditory event-related potential (ERP) [21]. In these experiments human subjects are exposed to a repetitive sequence of standard sounds, interrupted by a different deviant sound. Neuronal populations in the brain encode acoustic invariances specific to each speech sound, and a change in the sequence elicits a measurable signal, the MMN. This occurs whether or not the subject is paying attention. It can even occur in coma patients, in which case it is a predictor of a return to consciousness.

4 The Perception of Phonemes

Consider the perception of phonemes, and the role in language understanding of minimal pairs of words such as *kick* and *lick* or *ball* and *bell*: the change in a single consonant or vowel phoneme alters the meaning. Because of the human ability to distinguish phonemes they were used in the MMN experiments (mentioned above) as well as other sounds.

Interest in phonemes goes back a long way and Sapir's research in the 1930s showed the phoneme is "a cognitive construct that is so strong that it leads individuals to assert the existence of sounds that are not present, and deny the existence of sounds that are present" [29], quoted by [3]. The ability to distinguish between different phonemes is commonly called "categorical perception" in the field of child development, and its acquisition by very young infants is remarkable [28,15,6]. Recent neuroscientific work has identified locations of phoneme detectors in the brain, in both hemispheres [35].

However, it remains the case that phonemes can still only partially be detected by automated recognition systems. In very simple terms, speech recognizers typically have two components: an acoustic processor which extracts a list of candidate phonemes and a language model which compares possible choices of syllables or words (short sequences of phonemes) with stored examples in a very large data base of recorded speech. The output of the recognizer is based on a combination of probabilities derived from these information sources.

The probability of phonemes being correctly recognised by an automated system is related to a number of factors, for instance reading a prepared text gives better results than spontaneous speech. Whether the phoneme is a vowel or a consonant is relevant, as is its position in a syllable. A key factor is the saliency of the word in which the phoneme occurs: content words like nouns, verbs and adjectives are more likely to have canonical, consistent phonemic structure than function words. Research on an analysis of 4 hours of spontaneous telephone speech which was manually annotated by trained phoneticians reported that the word “and” had 80 different phonemic representations [7]. This highlights the point that orthographic transcripts of spontaneous speech may not be a close match to perceived sounds - a point that needs to be taken into account when research into, for instance, child language acquisition uses corpora of written transcripts of child directed speech. The observed fact that content words are more likely to have consistent phonemic structure is exploited in work simulating the learning of word forms by a humanoid robot which interacts with humans [17].

Problems with phoneme recognition are partially due to speakers’ pronunciation: human listners use a number of semantic, syntactic and prosodic clues to decode the speech signal. However, it remains the case that at present phonemes cannot reliably be detected by automated recognition systems, even when clearly enunciated. Recent research on phoneme recognition obtains best results of only around 67% [34,2]. This research, based on empirical evidence such as EEG brain wave recordings, suggests that phases of oscillations might account for phonemic discrimination. Port proposes a high-dimensional linguistic memory, incorporating many items of information extracted from an acoustic stream, where phonemes are just statistical invariants drawn from this data [27]. This type of interpretation has led him to question the existence of phonemes, since they are “only” cognitive constructs. There is no reliable external representation of this hearing experience, which contrasts with the traditional sensorimotor view in explaining visual consciousness that “the outside world is its own external representation” [25].

When phonemes are described as *only* cognitive constructs it implies that there could be some other, more real, status. For instance, with a robotic prosthesis we can say it is not real, it contrasts with a real limb. But this is not the case with phonemes: cognitive constructs are as real as you can get. “We must not, therefore, wonder whether we really perceive a world, we must instead say: the world is what we perceive” [32, page xvi].

5 The Varied Meanings of “Representation”

At this point we need to consider some of the varied meanings of the word “representation”. It can mean a symbol, as a flag can represent a country. It can mean a pictorial artefact, as in representational art, and an analogous use of the word in this sense was common in cognitivist theories of perception which proposed some sort of mental picture in the brain [32,25,13]. Representations might be thought of as “images, schemas, symbols, models, icons, sentences, maps and so on” [13].

However, the word is also commonly used rather differently to refer to neuronal patterns of activity in the brain. In the report on the drumming research mentioned above the authors write of areas of the brain that “are active in audiomotor, visuomotor and audiovisual *representation* studies” (emphasis added in this and following quotations). Wang [34] describes his work as investigating “the neural *representation* of phonemes”. Other examples of the use of the word include: “each sound, both speech and non-speech, develops its neural *representation* corresponding to the percept of this sound” [21]. Another example is: “comprehension is achieved with RH acoustic / phonetic *representations* of speech working in concert with LH mechanisms more sensitve to phonemic category” [35].

In this usage the term "representation" refers to a relationship beween neurons and external stimuli; there is no inner picture or reconstructed model. This contrasts with the use of "representation" to mean some mediating image beween an internal and external world.

With the perception of phonemes there is no identifiable external representation as with visual percepts, but there is a representation in the form of neuronal patterns of activity. Since the advent of brain imaging technologies there has been much research into the neuroscience of language and investigations into neuronal functions. For example, this approach addresses the relationship between speech perception and production, showing how they are critically linked, [10]. Another example reports evidence that phonological input and output buffers hold transient information, and a phonological short term' memory (pSTM) arises from the cycling of information between the two buffers [14]. Speech processing can be seen as processing statistical invariants extracted from an acoustic signal [27]. But though there is no identifiable external representation we still have sensorimotor interaction with a source.

6 Discussion

The language used in writings about perception often implicitly suggests that perception is typically visual. Thus talk of an "observer" perceiving "objects" or "things" implies vision, possibly also touch. We would not usually talk of observing sounds. However we need to recognize this bias towards vision, and take a more comprehensive view that includes all perceptual modes in our understanding of perception in both natural and artificial domains.

There is often a false dichotomy between competing theories of perception. Taking some of the most prominent features of a sensorimotor account we have looked at whether the perceiving subject is necessarily active. In vision, this is typically the case, but in other perceptual modes such as audition the subject can commonly be either active or passive. This does not mean that the theory of an "active" process in vision is mistaken: the mistake comes in claiming for all types of perception the necessity of a process that does not apply universally.

Another approach is to reconsider what we mean by "active". The train of thought from Merleau-Ponty and then Varela is that perception requires directed attention, or intentionality. In a contrasting neuroscientific approach the concept of "active" is viewed differently. For instance, Zeki comments on functional specialisation in visual processing - such that different areas of the brain process color, shape, movement autonmously. He says that this "has been instrumental in changing our minds about vision as a process, impelling us to consider it as an active process ....The brain, then, is no mere passive chronicler of the external physical reality, but an active participant ..." [36]. However, this activity in the brain is a subconscious process, not directed or intentional (in either sense). We can refer back to Heidegger's point that perception can become routine, or subconscious, as in turning a handle to open a door. We may only pay attention if there is some interruption to the usual routine. Now, this type of subconscious activity, in which the action can easily be restored to attention, is one end of a spectrum. At the other end we have neural activity, for instance in integrating the color and shape of an observed object, that we cannot control. In between there are a range of activities that can be moderated with training and effort - for instance the control of a robotic limb discussed above. After a stroke affecting a patient's motor abilities, conscious effort to control muscles that would normally function without thought can be part of a rehabilitation process.

Misconceptions about the sensorimotor account of perception can partly be traced to the ambiguity of the term "representation". It can be taken to mean a reconstructed mental image and a key part of the sensorimotor account is that there are no such images of objects perceived: their own external existence is their representation. However, the term "representation" can in contrast be used to refer to distributed neuronal patterns in the brain, and in this sense all perceptual modes are associated with a representation. Varela et al. discuss the different meanings of "representation" [32]. O'Regan and Noë write of cortical maps [25, page 939] which others might refer to as "representations". These distributed neural patterns are not reconstructed representations like a picture.

A core concept of a sensorimotor account of perception is that the perceiving subject is in the world, not separated by a mediating construction. Our understanding of what constitutes this world is much deeper than a collection of objects that can be seen or touched. It includes auditory perceptions and other perceptual modes found in non-human animals that could inspire robotic development [33]. By deepening our understanding of what constitutes the external world and how we interact with it does not diminish the sensorimotor account of perception, but gives it a firmer empirical base.

References

1. Arena, P., Patane, L.: Spatial Temporal Patterns for Action-Oriented Perception in Roving Robots. Springer (2009), doi:10.1007/978-3-88464-4
2. Baghai-Ravary, L.: Evidence for the strength of the relationship between automatic speech recognition and phoneme alignment performance. In: ICASSP (International Conference on Acoustics, Speech and Signal Processing) (2010)
3. Beeman, W.O.: Linguistics and Anthropology. In: Kempson, R., Fernando, T., Asher, N. (eds.) The Philosophy of Linguistics. Elsevier (2012)
4. Berkeley, G.: Essay towards a New Theory of Vision (1709)
5. Cangelosi, A., et al.: Integration of action and language knowledge: A roadmap for developmental robotics. IEEE Transactions on Autonomous Mental Development 2(3), 167–194 (2010)
6. Curtin, S., Hufnagle, D.: Speech perception. In: Bavin, E. (ed.) The Cambridge Handbook of Child Language. CUP (2009)
7. Greenberg, S.: Speaking in shorthand: A syllable-centric perspective for understanding pronunciation variation. Speech Communication 29, 159–176 (1999)
8. Hatsopoulos, N.G., Donoghue, J.P.: The science of neural interface systems. Annual Review of Neuroscience 249, 249–266 (2009)
9. Heidegger, M.: Being and Time. Blackwell (1927), translated by McQuarrie, J., Robinson, G. published (1962)
10. Hickok, G., Houde, J., Rong, F.: Sensorimotor integration in speech processing: Computational basis and neural organization. Neuron 69(3), 407–422 (2011)
11. Hume, D.: A Treatise of Human Nature. OUP (1740, 1978), book I, part IV, section II
12. Husserl, E.: Logical Investigations. Routledge and K Paul (1900), translated by Findlay, J. published (1970)

13. Hutto, D.: Radically enactive cognition within our grasp. In: Radman, Z. (ed.) The Hand: An Organ of the Mind. MIT Press (2013)
14. Jacquemot, C., Scott, S.K.: What is the relationship between phonological short term memory and speech processing? Trends in Cognitive Sciences 10(11) (2006)
15. Kuhl, P.: Early language acquisition: Cracking the speech code. Nature Reviews - Neuroscience 5, 831–843 (2004)
16. Locke, J.: An Essay Concerning Human Understanding. Thoemmes, published 2003 (1690), book II, chap IX, section 9
17. Lyon, C., Nehaniv, C.L., Saunders, J.: Interactive Language Learning by Robots: the transition from babbling to word forms. PLoS1 7(6) (2012), doi:10.1371/journal.pone.0038236
18. Magee, B.: The Great Philosophers. OUP (1987), reprinted (2009)
19. Merleau-Ponty, M.: Phenomenology of Perception. Routledge, New York (1945), translated by Colin Smith, reprinted (2005)
20. Merleau-Ponty, M.: Husserl at the Limits of Phenomenology. NW University Press (1959), edited by Lawlor, L., Bergo, B. (2002)
21. Näätänen, R.: The perception of speech sounds by the human brain as reflected by the mismatch negativity (MMN) and its magnetic equivalent (MMNm). Psychophysiology 38, 1–21 (2001)
22. Nehaniv, C.L., Förster, F., Saunders, J., Broz, F., Antonova, E., Kose, H., Lyon, C., Lehmann, H., Sato, Y., Dautenhahn, K.: Interaction and Experience in Enactive Intelligence and Humanoid Robotics. In: IEEE Symposium on Artificial Life (IEEE ALIFE). IEEE Symposium Series on Computational Intelligence, SSCI (2013)
23. Noë, A.: Action in Perception. MIT (2006)
24. Noë, A., Thompson, E.: Vision and Mind. MIT (2002)
25. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, 939–1031 (2001)
26. Petrini, K., Pollick, F., Dahl, S., McAleer, P., McKay, D., Rocchesso, D., Waadeland, C.H., Love, S., Avanzini, F., Puce, A.: Action expertise reduces brain activity for audiovisual matching actions: An fMRI study with expert drummers. NeuroImage 56, 1480–1492 (2011)
27. Port, R.: How are words stored in memory? Beyond phones and phonemes. New Ideas in Psychology 25, 143–170 (2007)
28. Saffran, J., Aslin, R., Newport, E.: Statistical learning by 8-month-old infants. Science 274, 1926–1928 (1996)
29. Sapir, E.: Le realite psychologique des phonemes. Journal De Psychologie Normale Et Pathologique, 247–265 (1933)
30. Saunders, J., Nehaniv, C.L., Lyon, C.: The acquisition of word semantics by a humanoid robot via interaction with a human tutor. In: Dautenhahn, K., Saunders, J. (eds.) New Frontiers in Human-Robot Interaction. John Benjamins (2011)
31. Tikanoff, V., Cangelosi, A., Metta, G.: Integration of speech and action in humanoid robots: iCub simulation experiments. IEEE Transactions on Autonomous Mental Development 3(1) (2011)
32. Varela, F., Thompson, E., Rosch, E.: The Embodied Mind. MIT Press (1991)
33. Vernon, D., von Hofsten, C., Fadiga, L.: A Roadmap for Cognitive Development in Humanoid Robots. Springer (2010)
34. Wang, R., Perreau-Guimaraes, M., Carvalhaes, C., Suppes, P.: Using phase to recognize English phonemes and their distinctive features in the brain. In: Proceedings of the National Academy of Science, vol. 109(50) (2012)
35. Wolmetz, M., Poeppel, D., Rapp, B.: What does the right hemisphere know about phoneme categories? Journal of Cognitive Neuroscience (2011)
36. Zeki, S.: Inner Vision: An exploration of Art and the Brain. OUP (1999)

From a Sensorimotor to a Sensorimotor++ Account of Embodied Conceptual Cognition

Joel Parthemore

Centre for Cognitive Semiotics, University of Lund, Sweden
joel.parthemore@semiotik.lu.se

Abstract. Since the publication of O'Regan and Noë's original article in *Behavioral and Brain Sciences* in 2001, which first set out the sensorimotor account by which sensory experience and motor engagement are inextricably intertwined, there have been not just one but many sensorimotor accounts. However, in many ways that original article remains the canonical account. In this paper, I discuss a particular theory of concepts from philosophy of mind – the *unified conceptual space theory,* based on Peter Gärdenfors' *conceptual spaces theory* – and, in that light, set out what I take to be the key points of the 2001 account, along with its strengths and weaknesses. I discuss the ways in which the 2001 account aligns with, and departs from, the *unified conceptual space theory*; and I offer an extension to it that I call *sensorimotor++*, which adds to the 2001 account a key role for emotional affect and the somatosensory system, with which one might ground salience, and a key role for (so-called 'mental') representation, properly understood. I argue that *sensorimotor++* makes for a better theory of concepts – one that is not just *embedded* and *embodied* but *enactive* – and, perhaps, a better sensorimotor theory more broadly.

Keywords: sensorimotor theory, concepts, representations, enaction, enactive, circular causality, conceptual spaces theory, unified conceptual space theory.

1 Introduction

Theories of concepts represent a sub-domain of philosophy of mind with considerable overlap into psychology and cognitive science, represented most prominently by such contemporary writers as Jesse Prinz [43], Jerry Fodor [10], and Peter Gärdenfors [13], with input from e.g. Ruth Millikan [28] and Alva Noë [31], and with a healthy criticism of the whole endeavour from e.g. Edouard Machery [24]. Concepts may be understood as the building blocks of systematically, productively, compositionally, and endogenously controlled structured thought, while conceptual abilities are those skills by which certain agents we identify as conceptual agents are able to cognize in a systematically, productively, etc. – and, above all, flexibly – structured fashion. I trace the notions of *systematicity* (the same concepts can be used in more or less the same fashion across unboundedly

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_10,

many contexts) and *productivity* (a finite number of concepts can be combined to form an unbounded number of complex concepts) to Gareth Evans' *Generality Constraint*, set out in [8, 100-104]. *Compositionality* (the ability of concepts to be joined together or taken apart) follows directly from the first two properties. I owe the phrase 'endogenously controlled' to Prinz (see e.g. [43, p. 197], who offers it as an alternative to the potentially misleading '(Kantian) spontaneity': concepts are not just passively given to the conceptual agent but somehow actively under her control. Note that nothing in this list entails that conceptual agents necessarily possess (human-style) language.

One may approach concepts and conceptual cognition in one of two ways, which I take to be equivalent. One talks of concepts as reified entities: the aforementioned "building blocks" of structured thought. The other talks of the abilities by which certain agents are able to engage with the world in cognitively creative ways. They are, to me, as two sides of a coin.

Meanwhile, the *sensorimotor*[1] *account* [34] offers a theory about the nature of cognition more broadly, where 'cognition' may be understood in rough-and-ready terms as the encounter of cognitive mind with physical world (in a way that rejects any kind of Cartesian substance dualism – more on that below); or, in philosophical terms, the means by which certain agents effectively create an online/offline distinction in their interactions with their environment: there is the world, and then there are thoughts about the world.

1.1 Toggling between Perspectives

Within the tea cup that is the field of concept studies, many a storm has brewed over whether concepts are (or are best understood) as abstract (objects) or concrete (abilities), representational or non-representational, public or private, atomic or structured (see Figure 1): so e.g. on Jerry Fodor's *informational atomism* account [10], concepts are atomic, public, representational, and abstract; on Jesse Prinz's *proxytypes* account [43] or Peter Gärdenfors' *conceptual spaces* account [13], they are structured, both public and private, representational, and abstract; on Noë's sensorimotor-theory-based account [31], I believe they are best understood as structured, public, non-representational, and concrete.

The problem is not, I think, that the concept of concepts is polysemous, as many might claim. The problem is more that we are standing too close to what we are trying to examine. On my own, *unified conceptual space* account [39,36,40], which is based on Gärdenfors' work, concepts are *either* abstract (objects) or concrete (abilities), *either* representational or non-representational, *either* structured or atomic, depending on which of two perspectives – both necessary to any proper theory of concepts – one is taking at any particular time. If, most of the time, conceptual agents must, logically, get on with possessing and employing concepts without stopping to think about their concepts *as* concepts (the one

[1] The term 'sensorimotor' is meant to capture that which is necessarily and simultaneously sensorial and motor-based: senses and motor system are not two separate things to be brought together but two sides of one coin.

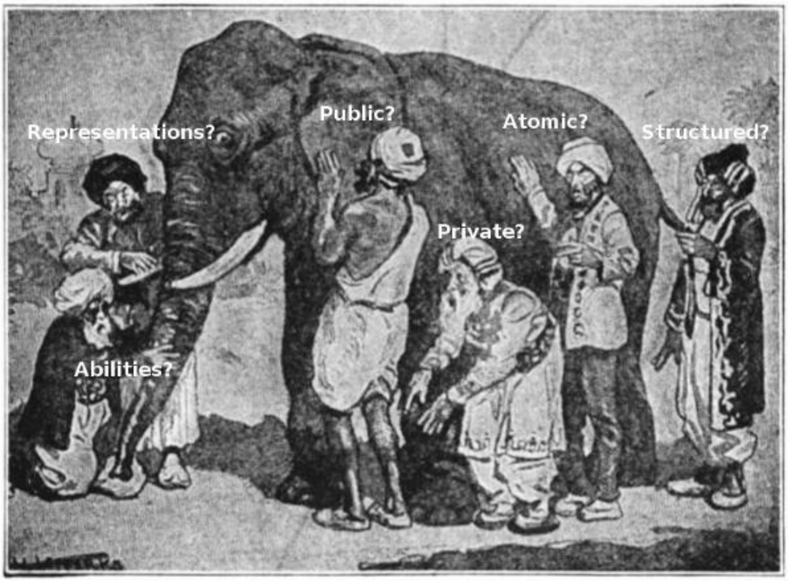

Fig. 1. The Conceptually Blind Men and the Elephant (picture downloaded from Wikimedia Commons: http://commons.wikimedia.org/ and edited), as originally appeared in [36]

perspective); then sometimes – for certain conceptual agents at least – the concepts *themselves* become the focus of attention (the other perspective). I believe that self-reflective conceptual agents toggle between these two perspectives constantly and, for the most part, non-reflectively: i.e., they do not stop to reflect on their own reflection (for practical reasons, as much as any fear of infinite regress!). Both perspectives are required, for anything like a complete view; but they cannot be resolved into a single perspective, on pain of contribution. Doing so would require setting our conceptual nature aside and stepping outside the observational system of which we are, inextricably, a part.

At the same time, on the *unified conceptual space* account, concepts are both public and private: again, depending on what perspective one is taking. One can talk about concepts for any given conceptual agent; for conceptual agents who are socially organized, one can talk about shared concepts for the group, which both relate closely to and yet differ from the individual agents' concepts, as described very nicely by Prinz [43, p. 159]; for conceptual agents who possess human-like language, one can talk yet again about the words of a language by which the shared concepts are (perhaps imperfectly) lexicalized.

1.2 Escaping the Cognitivist Trap

For all of the frequent talk by Noë, by Vittorio Gallese and George Lakoff [12], and by many others of concepts as concrete abilities or direct engagements with the world, concepts still often come off sounding highly abstract and far removed from the world. The target of complaint is variously described as cognitivism, symbolic AI, GOFAI ('good old-fashioned AI') – a term coined by John Haugeland in [18], or Cartesian mind/body dualism. The concern is that cognition comes off as disembodied and that the 'offline' mode is understood as actual detachment from – rather than inattention to – the environment. It creates an unintended and unnecessary *explanatory gap* between mind and body, cognitive agent and physical world: how to bridge the divide?

The peculiar, regress-inviting way by which we (necessarily, I think) conceptualize about our concepts should not – these researchers would say – mislead us into over-intellectualizing their nature: concepts ultimately are not about intellectual 'exercise' but hands-on activity. In keeping with them and *contra* e.g. Fodor's *informational atomism*, I want to reject the notion of concepts as physical symbol systems [30]: one stage in an input-output-based, *SMPA (sense-motivate-plan-act)* architecture easily describable in terms of a software program running on a digital computer. As Noë writes [31, p. 2]:

> . . . We ought to reject the idea – widespread in both philosophy and science – that perception is a process in the brain whereby the perceptual system constructs an internal representation of the world. No doubt perception depends on what takes place in the brain, and very likely there are internal representations in the brain (e.g., content-bearing internal states). What perception is, however, is not a process in the brain, but a kind of skillful activity on the part of the animal as a whole.

Remember what I wrote earlier about concepts being one thing when we stop and reflect on them but logically another when we get on with possessing and employing them non-reflectively. The trick is that, the moment we reflect on what it is to possess and employ them non-reflectively, we bring them within the domain of (seemingly internal) representations! That does not mean, however, that we are powerless to say anything about what possessing and employing them non-reflectively could mean; and here something like sensorimotor theory, grounded in the empiricist tradition, seems to me to offer the best way forward, offering a way to trick ourselves out of the cognitivist trap.

1.3 The *Unified Conceptual Space* Theory

Space prevents me from recapitulating the *unified conceptual space theory* here (see the above references), other than to offer a very brief summary. In brief, however, conceptual spaces theory, on which the unified conceptual space theory is built, is a prototype- or similarity-space-based theory of concepts couched in the language of geometry, whereby concepts arise and evolve by the progressive partitioning and re-partitioning of conceptual spaces described as Voronoi tessellations (see Figure 2). The theory of prototypes – best examples of a category – derives from the work of Eleanor Rosch [45,44]; while a *similarity space* locates concepts in an abstract space defined by a set of integral dimensions with a predefined metric, such that the closer two concepts (points or sub-regions) are located within the space, the more similar they are judged to be.

The unified conceptual space theory attempts to fill in some of the missing detail in conceptual spaces theory, whilst pushing it in a more algorithmically amenable, more empirically explorable way, inspired by the work in prototype theories in general and *conceptual spaces theory* in particular. It provides a specific algorithm – a kind of recipe – by which one can move from protoconcepts (foundational elements that fail to meet one or more of the properties offered above) to concepts to concepts of concepts. That algorithm is currently best described as semi-formal, though clearly detailed enough to allow direct translation of the theory into a mind-mapping software program (see Figure 2).[2] The intention is to develop the algorithm into a properly formal (non-monotonic) logic, albeit one that allows for the apparent inconsistencies in most people's conceptual frameworks. The thinking is that inconsistencies at the global level can be tolerated, provided they are sufficiently spatially removed from each other in the unified space (e.g., thinking X at Time T and *not* X at Time T+100); what cannot be tolerated is localized inconsistencies (e.g., thinking X and *not* X at *Time T*).[3]

The recipe looks like this: for any given conceptual agent, concepts (of whatever type) are located within a common *space of spaces* that brings together all the many conceptual spaces described by *conceptual spaces theory.* An analogous space is meant to exist at the group or societal level, as inspired by [22] and, even more so, [49].

Concepts have both proximal and distal connections to one another. The proximal connections are along the three dimensions that define the unified space: what I call the *axis of generalization* (the familiar concept hierarchy: a dog is a mammal is an animal. . .), the *axis of abstraction* (from 'lower order' / concrete / 'physical' to 'higher order' / abstract / 'mental'), and the *axis of alternatives* (obtained by varying the values of any one or more integral dimension according to a pre-given metric: e.g., *colour* has the integral dimensions of *hue, saturation,*

---

[2] That such a translation is possible should not be surprising: the theory was designed, from the beginning, with such an application in mind. For more on mind-mapping programs, which are intended to allow users to brainstorm ideas and to 'externalize' their understanding of one or another conceptual domain, see e.g. [33] and [47, pp. 77-82].

[3] For some readers, the worry will persist that *any* logic-based treatment of concepts or conceptual cognition will both over-intellectualize matters and fail to capture such a dynamic view of cognition as the enactive approach (see below) is committed to. The implied claim – that an appropriately designed algorithm or logic *can* capture such a dynamic view of cognition – is beyond the remit of this paper. Here, I will simply note Rick Grush and Patricia Churchland's response to Roger Penrose, where they point out that it remains very much an open question whether there is *any* phenomenon in the physical universe that cannot, at least in principle, be algorithmically described (p. [15, p. 190]). I readily allow that, in practice, an algorithm for concepts is likely – indeed, almost certain – to leave much out; but the *unified conceptual space theory* makes no claim to be the final word on concepts.

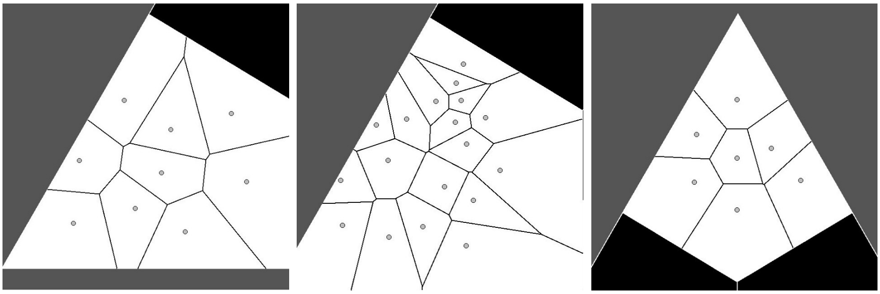

Fig. 2. A sample of Voronoi tessellations, created by the mind-mapping program described in Chapter 8 of [36] as a direct translation of the *unified conceptual space theory* into software

and *brightness*). All three dimensions are divergent in both directions, meaning that the geometry of the space is hyperbolic rather than Euclidean. Meanwhile, distal connections are of three types: some concepts describe a *component* (mereological) relation to other concepts; some describe *parameters* or *properties* to other concepts; and all may be *contextually related to* (commonly associated with, but not required by) other concepts.

In Section Two, I set out what I take to be the pluses and minuses of sensorimotor theory, with respect to theories of concepts in general and the *unified conceptual space theory* in particular. Section Three proceeds to describe in greater detail how the canonical sensorimotor account and the *unified conceptual space theory* view cognition differently. Section Four presents *sensorimotor++*: an extended version of the sensorimotor account formulated as part of the *unified conceptual space theory.* Section Five summarizes the discussion and offers some conclusions.

2 Sensorimotor Theory: Pluses and Minuses

Before I discuss how sensorimotor theory, as described in [34], may be adapted to work with the *unified conceptual space theory*, I first must address how it does – and does not – fit that theory.

2.1 Pluses

The key points of sensorimotor theory, as described in [34] and implied by the quote above, I take to be these, all of which are strongly endorsed by the *unified conceptual space* account:

– All mental content, conceptual or otherwise, must be grounded in specific sensorimotor engagements, making sensorimotor engagements at least partly constitutive of that content.[4] This places sensorimotor theory squarely in the empiricist tradition embraced by *proxytypes theory*, *conceptual spaces theory,* and the *unified conceptual space theory* and in contrast to the rationalist tradition embraced by *informational atomism.*[5]

– Mental content is not fixed but dynamic and, indeed, contingent: *if* I do this, *then* I will experience that: a position that has much in common with so-called *ecological psychology* [14]. This accords as well with the enactive philosophy implicit in *conceptual spaces theory* (as I have described e.g. in [39]) and explicit in the *unified conceptual space theory.* Associated with Francisco Varela, Humberto Maturana, and Evan Thompson among others (see e.g. [27,26,54,51]), the spirit of enactivism is perhaps best summarized in a poem of Anotonio Machado, translated by Varela [53, p. 63]):

  Wanderer, the road is your
  footsteps, nothing else; wanderer, there is no path,
  you lay down a path in walking.
  In walking you lay down a path...

– In keeping with this, mental content is part-and-parcel with interaction: input is logically inseparable from output, except in certain – very limited and circumscribed – cases that should not be mistaken for the general case. This, too, accords with enactivism, by which neither pre-conceptual agent nor pre-conceptual environment are recoverable from their interaction: 'I have proposed using the term *enactive* to... evoke the idea that what is known is brought forth, in contrast to the more classical views of either cognitivism or connectionism' [27, p. 255]. Indeed, Noë for a time called his version of sensorimotor theory, developed in the years following the 2001 paper with O'Regan, 'enactive' (see e.g. [31]), though in more recent years he has preferred to talk of *actionism* (see e.g. [32]).

O'Regan and Noë introduce the key phrase *sensorimotor contingencies* to

[4] Contrast this with e.g. the position described in [12], where all mental content, including the most abstract, *just is* specific sensorimotor engagements, albeit with parts of those engagements routinely suppressed (when one thinks or talks of democracy, one need not 'do' democracy): that is to say, there is no residue to be explained once sensorimotor engagements have been accounted for. The account is oddly reminiscent of George Berkeley's discussion [1] of triangles and his argument that no one has an abstract concept of triangle that is anything more or other than a specific triangle instance (for Berkeley, a mental picture of a triangle). Like Berkeley, Gallese and Lakoff reject the notion of abstract classes in favour of 'concrete' instances.

[5] Sensorimotor theory constitutes a refinement and renewal of classical empiricism (I have in mind in particular the British empiricism of John Locke, David Hume, and George Berkeley) – as Prinz has tried to do in his own way with his proxytypes theory. *Contra* rationalism, for all of the importance of reasoning to (human) cognition, cognition is ultimately grounded not in reason but experience. For all that reasoning may conclude that something surely must logically be impossible, experience can always come along and say: no, in fact; here is the proof that it cannot be!

describe what governs these interactions. Here is how they describe holding a bottle [34, p. 945]:

> In fact... you may well have very *little* sensory stimulation coming from the bottle at the present moment. Yet, you actually have the feeling of 'having a bottle in your hand' at this moment. This is because your brain is 'tuned' to certain *potentialities*: if you were to slide your hand very slightly, a change would come about in the incoming sensory signals which is typical of the change associated with the smooth, sliding surface of glass.

- Most of cognition has little to do with representation, in any but the most loosely metaphorical of senses. Whatever representation there is and whatever representation is about, 'there is no "*re*"-presentation of the world inside the brain' [34, p. 940]. This is because, most of the time, cognition is not something we think about: it's something we *do,* an 'exploratory activity' [34, p. 940] where one experiences the world by exploring it. What one cannot actively explore, one cannot experience (see e.g. [19]). Remember what I said earlier that, most of the time, conceptual agents just get on with possessing and employing their concepts non-reflectively; it is only when they stop to examine their concepts that their concepts become representational.
- Sensorimotor theory recognizes the important difference between so-called mental imagery and pictures – a point that Göran Sonesson has long argued (for a recent treatment, see [48]). According to the classical imagist account of concepts, as described in [43, pp. 25-31], concepts just *are* conscious pictures in the mind. He writes [43, p. 139]: 'to bring concept empiricism up to date, one must abandon the view that concepts are conscious pictures'. Instead, for Prinz, concepts are partly unconscious multimodal 'images'. According to the *unified conceptual space theory*, conflating mental imagery with pictures (conscious or otherwise) confuses the very nature of what they have in common: namely, what makes them both (under certain circumstances) representations (see Section 3.3 and [38]). Unlike pictures, concepts are not primarily things that we 'see', even metaphorically.

In sum, sensorimotor theory consciously tries to avoid over-intellectualizing cognition: a surely laudable effort. If we think too much about how we think, we risk losing our very target in an elaborate set of mind games (not to mention, we will never be seen to *do* anything, because we will always be thinking about *what* to do). Sensorimotor theory takes on board the best of Rodney Brooks and the lesson that, if an agent can offload some amount of its cognition onto the world – knowing that the world will be one way and not another (the source of the – in my mind – misleading and ill-advised proverb that 'the world is its own best model' [2, p. 5][6]) – surely it will. Putting this another way, sensorimotor theory reduces the number of steps needed for explaining a number of key cognitive abilities. If you don't need to build more in to make it work, don't: a fine

[6] *Cf.* [34, p. 939]: 'the outside world serves as its own, external, representation'.

application of Occam's razour. The lesson is one that theories of concepts – for all their inherent tendency to over-intellectualize cognition – would do well to mind.

2.2 Minuses

As described in the original 2001 paper, sensorimotor theory assumes – indeed, depends on – a particular, realist metaphysical position without making that commitment clear: a charge I would similarly lay on Fodor's *informational atomism* theory, which stands or falls on its brand of realism. If one does not share that position, then one might not feel that the so-called *hard problem of consciousness* [5] – which the 2001 paper boldly purports to dissolve – has really been addressed. Indeed, both *conceptual spaces theory* and the *unified conceptual space theory* are inclined toward – if not necessarily dependent on – an antirealist metaphysics (a position that, indeed, much of the enactive community, with its emphasis on the inseparability of agent and environment, tends toward as well). At one point, the 2001 paper does come close to acknowledging that its position is more metaphysical than empirical [34, p. 948]: 'we are providing a general framework for the study of vision, and it is not possible to subject a general framework to direct verification. Our new framework provides scientists with new problems and it makes some old problems appear as non-problems. . . '.

Both the 2001 paper and, even more so, Noë's subsequent book [31], make strong assumptions about what people experience that one need not necessarily share: e.g., 'many people say that they have the impression that when they see, the entire visual field is present to consciousness in all its nearly infinite detail' [34, p. 961]. Perhaps 'many' people do; but many, I would claim, do not. Similarly, Noë writes [31, p. 56]: 'if you reflect on what it is like for you to look at the wall, you will notice that it seems to you as if the whole wall is there, at once. . . ' and [31, p. 57] 'of course, it *does* seem to us as if we have perceptual access to a world that is richly detailed, complete, and gap-free. And indeed we do! We take ourselves to be confronted with and embedded in a high-resolution environment'. This strikes me as very much a philosopher's and not a lay person's way of putting things: the very error Noë is so anxious to avoid. As a philosopher, I am frequently and painfully made aware that non-philosophers (or other philosophers for that matter) do not necessarily view the world as I do. In any case, a theory of concepts – if it lays *any* claim to being a general theory of concepts – must be able to capture that portion of people's experience that is conceptually structured[7] in all its richness *or* poverty.

Although paying lip service to other sensory modalities – with the sensorimotor account is also meant to explain – the 2001 paper focuses much if not most

[7] I leave aside, for now, the question of whether experience is *entirely* conceptually structured, as the so-called conceptualists would hold, or a mixture of the conceptual and non-conceptual, as the non-conceptualists contend – other than to say that, on this point, I side with the non-conceptualists. What matters here is that both sides of the debate agree that *at least part* of experience is conceptually structured.

of its attention on *visual* cognition and experience. To the extent that one is not attempting to explain all of cognition and experience, this is appropriate: one explains a part in the hope that it might be useful to the explanation of other parts. To the extent though that one is privileging vision above the other sensory modalities – as, I believe, many philosophers are implicitly inclined to do – one may well be, to pardon the metaphor, distorting the picture. Certainly in the context of conceptual cognition, too narrow of a focus on vision (or indeed any one modality) is not helpful, as the earlier discussion about classical imagism should have made clear.

O'Regan and Noë make a primary-vs.-secondary quality distinction – echoing John Locke – of which I am quite skeptical, given the spectre it raises of mind/body dualism. Primary qualities, of course, are meant to be 'in' an entity, independent of observer or context, whereas secondary qualities depend on observer and context. O'Regan and Noë apply this distinction not only to the objects of visual attention but to vision itself. There are, for example – at least in the 2001 paper – two kinds of sensorimotor contingencies [34, p. 943]:

> Sensorimotor contingencies of the first sort – those that are determined by the character of the visual apparatus itself – are independent of any categorization or interpretation of objects and can be considered a fundamental, underlying aspect of visual *sensation*. Sensorimotor contingencies of the second sort – those pertaining to visual attributes – are the basis of visual *perception*.

I am inclined to push the earlier point about interaction further and suggest that both primary and secondary qualities, and both kinds of sensorimotor contingencies, arise from the interaction of agent and environment in such a way that observer and observed cannot cleanly be disentangled: i.e., the interaction is fundamental or foundational.[8] This is, I think, what separates sensorimotor theory from enactivism – or, if you will, from other enactive accounts. This is how enactivism incorporates, but ultimately goes beyond, notions of embeddedness or *situatedness* (the cognitive agent is always located in a specific environment, which shapes its cognition) and embodiment (the cognitive agent always takes a particular physical form, which likewise substantively shapes its cognition). With the 'secondary' qualities and the one kind of sensorimotor contingencies, the role of the observer is explicit and unavoidable; with the 'primary' qualities and the other kind of sensorimotor contingencies, the role of the observer is implicit: the observer sits in the background, out of sight. Sometimes, attention will be more on the cognitive agent to the exclusion of physical environment, making one set most prominent; other times, attention will be more on the physical environment to the exclusion of the agent, and the other set will be emphasized. Either way, the observer is always present and must be accounted for: that is, both 'primary' and 'secondary' qualities, and both kinds of sensorimotor contingencies, depend

[8] Note that this is really a metaphysical claim rather than one argued from the evidence: that is, it is taken as a starting assumption, whose merits or lack thereof lie in what conclusions it may lead to: *if* one assumes this, *then* this is what follows.

on the observer – without whom, there are neither qualities nor contingencies. As Humberto Maturana writes, '*everything that is said, is said by an observer to another observer that could be himself*' [25, p. 30]. Inman Harvey writes, 'the underlying assumption of many is that a real world exists independently of any observer; and that symbols are entities that can 'stand for' objects in this real world in some abstract and absolute sense. In practice, the role of the observer in the act of representing something is ignored' [17, p. 5].

O'Regan and Noë are at pains to suggest that sensorimotor theory – unlike, presumably, its competitors – avoids any recourse to magic: e.g., [34, p. 946] '...by taking the stance that the experience of vision is actually *constituted* by a mode of exploring the environment, we escape having to postulate magical mechanisms to instill experience into the brain'. Avoiding magic is fair enough; but what exactly *is* magic? The supernatural, like the natural, is often discussed and rarely defined – as if everyone just knew already what it means. I am inclined to borrow a page from the science fiction author Arthur C. Clarke and suggest that magic is that which we do not currently understand – with the caveat that some things may not be possible for us to understand, even in principle, because they lie beyond our conceptual abilities (such as e.g. what it means to imagine a tesseract in all its four-dimensional glory). If *that* is what magic is, then, while I am inclined to agree with O'Regan and Noë that a sensorimotor-based approach leaves *less* explanatory residue, that is not to say there will be none remaining. If, by a 'fully naturalized account', one means an account that is both complete and consistent, then I would borrow a page from Douglas Hofstadter [20] in suggesting that, for any sufficiently expressive system, completeness and consistency rarely if ever sit comfortably together: one should *expect* explanatory residue (and not only when describing tesseracts!). Indeed, in [36], I argue – as part of the *unified conceptual space theory* – that, *contra* Roger Penrose [41, esp. pp. 72-77], there are good reasons to believe that conceptual understanding in general is knowably bounded even while there is no reason to think that the mind-independent world is similarly bounded.

Finally, for all of the importance of the 2001 paper and ensuing discussions, the sensorimotor account was never quite as new as it presented and presents itself. Its biggest contribution, I think, lies not in what it brings new to the table but in its way of describing things: things that many of us at least part way understood but lacked the adequate language for.

3 From Sensorimotor toward Sensorimotor++[9]

So far I have attempted to describe how the *unified conceptual space theory* aligns with, and where it departs from, the original sensorimotor account. Before I describe what the *sensorimotor++* account, first described in [40] and again in [36], adds to to that account – i.e., what each of the pluses stands for – I need

[9] The name 'sensorimotor++' was suggested to me by Peter Gärdenfors following a discussion with him about the draft of the paper that would become [40], later incorporated into Chapter 7 of [36].

to address in more detail how its outlook on cognition (and, by extension, that of the *unified conceptual space theory* of which it forms a part) differs.

3.1 Causality

> *Sensorimotor theory favours a linearly structured account; sensorimotor++ opts for a circular causality.*

On the canonical account, sensorimotor engagements give rise to sensorimotor profiles and sensorimotor profiles to so-called higher cognition: a largely if not strictly uni-directional, bottom-up process from the sensorimotorly concrete to the conceptually abstract, from mechanically driven associations to flexible conceptual structures. This may be a consequence (see Section 3.2) of sensorimotor theory's often very strong externalism and the consequent primacy it gives (on my reading) to the environment driving the agent rather than the agent driving the environment.

Clearly – as I am using the terms – not all cognitive agents are conceptual agents: that is, regardless of whether cognition 'goes all the way down' to the simplest organisms, as enactivism tends to favour (see e.g. [50]), conceptual cognition does not.[10] In the case of cognitive agents who are *not* conceptual agents – which I take to be a majority – the bottom-up linear-causal account may well be the most appropriate. Such agents will be like the purely stimulus-response-driven automata Descartes envisioned all non-human animals to be.

Nevertheless, in the case of conceptual agents – those that meet all the desiderata offered in Section 1 – it seems to me that one can equally turn the perspective around, to look at how sensorimotor engagements consist of or are built upon *conceptually structured* perceptions: how mind constrains the (experienced) world. Experience gives rise to concepts, which, in turn, structure experience; it is logically impossible, as a conceptual agent, to set one's concepts aside and step outside of the loop, to see the world 'as it really is': to do so would be like the dragon swallowing its own tail. One can at most gesture at what the world 'outside the loop' must be like. As I wrote in [40, p. 297]:

> Concept acquisition and application go hand in hand. Acquiring concepts is a process of applying concepts, which may themselves change in the process of acquiring new concepts. . . . Our conceptual spaces, individually and collectively, are both the product of our interaction with our environment and the basis for it. The model of causality is not linear but circular.

[10] Indeed, I am more inclined to the sort of account of 'how far down cognition goes' offered by Andy Clark and Rick Grush [6,16], where cognitive agents are those who, at least to a limited extent, are able to step back from the here-and-now and create a functional distinction to the 'was' or 'might be'. I am indebted to one of the anonymous reviewers for suggesting the clarification and the references.

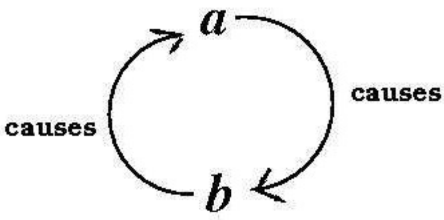

Fig. 3. A model of circular causality: *a* causes *b* causes *a*

Note that, a circular causal account (see Figure 3) is only coherent if the two 'sides' of the circle are considered independently of one another: i.e., as two separate instances of linear causality. Considered jointly, *no consistent interpretation can be given.*

3.2 Externalism

> *Sensorimotor theory is often strongly externalist; in keeping with its enactive perspective, sensorimotor++ seeks to avoid what it sees as the equally undesirable extremes – the Scylla and Charybdis if you will – of externalism and internalism.*

Remember that, for the enactive perspective even far more than for sensorimotor theory, everything keeps coming back to interaction: both experienced agent and experienced environment are a product of that interaction; agent does not separate cleanly from environment nor mind from world. On an enactive account, cognition in general and concepts in particular are not in the (brain of) the agent (*internalism*) nor in the agent's environment (various forms of *externalism*); insofar as they can be located anywhere, they are metaphorically 'in' the interaction of agent with environment. It is in the nature of concepts always to be setting boundaries and creating categories; it is in the nature of the enactive perspective always to view those boundaries as flexible and, indeed, dynamic: conceptually drawn lines that mask underlying continuities. As the author has heard one prominent researcher in the field to say, enactivism does not think much of fixed boundaries.

In keeping with the enactive tradition, when talking about cognition or cognitive phenomena, one must, I think, take great care using terms like 'inside' and 'outside'. As I argue in [37], these are terms that apply to physical volumes, and their usage in any other contexts should be understood as loosely metaphorical at best: neither mind nor cognition nor sensory perception is *prima facie* a physical quantity even as they are all physically realized.

The result of sensorimotor theory's externalism is, I believe, a tendency to overlook or downplay the role of the agent's bodily states as interoceptively and proprioceptively experienced; and, in particular, to overlook or downplay the agent's emotions.[11] Indeed, this is the diagnosis offered by Anthony Morse and Tom Ziemke in an unpublished paper from 2010, where they call their proposal the *somatic sensory hypothesis* [29]. As they argue there, without an account of emotions and consequent motivations, sensorimotor theory cannot ground salience: it particular, it cannot explain why some affordances are salient and others not, nor why and how those saliences change over time. Although many philosophers, including Aaron Sloman[12] see rationality and emotion as opposed to each other, the *unified conceptual space theory* is inclined to see emotional affect as part of the necessary foundation to cognition in general – and, therefore, to conceptual cognition in particular.

3.3 Representationalism

> *Sensorimotor theory is often read as moderately or even strongly anti-representationalist; sensorimotor++ favours a qualified (or 'modest') representationalism.*

Noë [31], in particular, I take to be anti-representationalist, in a way not so far removed from Brooks [3] or John Perry [42]. An informal poll would suggest that many if not most enactive philosophers are strongly anti-representationalist as well: consider e.g. the *radical enactivism* advocated by Daniel Hutto [21], for whom there is no such thing as mental content, let alone representation.

I share with O'Regan and Noë their distaste for so-called internal representations and the pride of place so often and uncritically, in certain circles – notably the cognitivist and GOFAI circles noted earlier – given to them. Indeed, in [38] I question the very coherence of talking about internal or mental representations as something ontologically distinct from any other kind of representation. Instead, I take the cue from Harvey [17] in restricting representation to a four-place function whereby agent Q uses R to stand in for S for agent T (who could, in certain circumstances, be agent Q herself). In [38] I take representation to be not a thing – let alone a reification – but rather an intentionally reflective perspective that certain agents, in certain circumstances, take toward one aspect or another of their experience, be it a picture or painting, or a thought in the mind.[13] As such – and in keeping with any perspective one takes – they are neither 'internal' nor 'external'. Like 'inside' and 'outside', these terms apply most appropriately to physical volumes and only in a loose metaphorical sense to anything else.

[11] It should be said that O'Regan acknowledges (personal communication) that the lack of any somatosensory account is one of the significant limitations of the 2001 paper – one that his notion of 'bodiliness' (see e.g. [35]) goes some way toward addressing; see also [4] for one suggestion of a formalized way to better take bodily states into account.

[12] Email correspondence.

[13] In this way my definition is, indeed, more restrictive than Harvey's, since Harvey does *not* require any degree of intentionality and so is willing to attribute 'minimal representations' to his artificial-life creations.

Further, as noted in Section 2.1 and implied at several points through this paper, I share with O'Regan and Noë – and, indeed, Brooks and Perry and others – their belief that much if not most cognition can be explained without resort to representations anywhere *except in the eye of the observer*. Nevertheless, as I said in the introduction, when it comes to concepts and conceptual cognition it seems that representations and representational language are difficult if not impossible to avoid. If, as is doubtless true, most of the time we possess and employ our concepts without reflecting on them *as* concepts; nevertheless, when we do, representations are what we find: concepts do not simply *re-present* the world (as a photograph is often naively thought to do) but *represent* it as being one way and not another, with the possibility that the 'picture' they describe could be (according to intersubjective consensus at least) wrong.

3.4 Empiricism vs. Rationalism

> *Calling to mind Gilbert Ryle's classic distinction* [46], *sensorimotor theory focuses on* knowing how *to the exclusion of* knowing that. *Sensorimotor++ suggests that, when it comes to concepts, concepts are neither precisely* knowledge how *nor* knowledge that *but something of neither and both. Most properly: concepts sit between* knowledge how *and* knowledge that.

The longstanding debate in philosophy between rationalism and empiricism can be understood as a disagreement over which is ultimately foundational to (what we recognize as) cognition: reason grounded in *knowledge that* (the sort of stuff one can reflect upon or *represent* to oneself and others) or perceptual discovery grounded in *knowledge how* (the sort of stuff one can generally reflect upon only poorly and represent or describe inadequately: consider the difference between knowing how to ride a bike and knowing precisely *what* that knowledge consists of).

In keeping with sensorimotor theory, *sensorimotor++* agrees that, when it comes to cognition in general, empiricism trumps rationalism, and the rationalist tendency to over-intellectualize should be strongly resisted. At the same time, *sensorimotor++* – and the unified conceptual space theory of which it forms a part[40,36] – sees conceptual cognition[14] as a special case: one in which *both* perspectives are needed and *neither* can be held as primary or ultimately correct. Concepts are both something that conceptual agents *do* and things they *possess*. They are both – as I suggested in the introduction – building blocks and abilities, and settling on one to the exclusion of the other is to miss not just half the story but – in a very real way – the whole thing.

This focus on processes over abstract objects, fluid and dynamic descriptions over reifications, is clearly evident in many places through the 2001 paper, but particularly in O'Regan and Noë's description of qualia [34, p. 960]:

[14] Remember that, on the *unified conceptual space theory,* conceptual cognition is taken to be only one small part of cognition.

> Qualia are meant to be properties of experiential states or events. But experiences, we have argued, are not states. They are ways of acting. They are things we do. There is no introspectively available property determining the character of one's experiential states, for there are no such states. . . . Experience is something we do and its qualitative features are aspects of this activity.

The problem with qualia, according to *sensorimotor++* and the *unified conceptual space* account, is that all of experience is meant to be structured from individuable quales. Instead, as described in [40,36], all experience – for the conceptual agent – is better understood as a mixture of conceptual and non-conceptual content (see Footnote 7); and while the conceptual content may be – indeed, probably is – appropriately structured, the non-conceptual content need not be.

4 The *Sensorimotor++* Account

Beyond these differences of emphasis and perspective, the *sensorimotor++* account, as one component of the *unified conceptual space theory*, adds two key ingredients to the 2001 version of the sensorimotor account: namely, an account of salience centered on the agent (albeit highly preliminary) and – for conceptual agents – a representationally disposed cognitive mechanism to actualize it. *Sensorimotor++* is sensorimotor engagements *plus* somatic and other bodily information (*per* Damasio [7] and Morse and Ziemke [29]) *plus* (appropriately qualified) representational language, as situated within a conceptual-spaces-based framework.

4.1 From Meaninglessness to Salience

Like the first airplane designers attempting flight, sensorimotor cognition faces a fundamental difficulty getting itself off the ground. Unless meaning is there from the very beginning, one is compelled to ask: where does meaning come from? Unless one believes that (even abstract) cognition consists in nothing more than specific sensorimotor engagements – as *per* e.g. Gallese and Lakoff [12] (see Section 4.2) – one faces a difficulty in how one moves *beyond* specific sensorimotor engagements: i.e., how one generalizes to the sensorimotor profiles one needs to explain affordances, nevermind abstract conceptual thought.

As said in Section 3.2, an improved sensorimotor account needs to give a key role to emotional affect and the somatosensory system more broadly. Without that, *sensorimotor++* argues, it cannot offer an adequate (or perhaps any) account of salience centred on the agent.[15] If the cognitive agent is not somehow predisposed to find certain aspects of the environment salient and others not, there will be nothing to guide or structure its sensorimotor engagements. If the

[15] One could argue, of course, that it is not the role of a sensorimotor account to explain salience, that that should be left to another theory.

conceptual agent is not somehow predisposed to seek out certain patterns in its environment and disregard others, there will be nothing to guide its developing conceptual frameworks down one or another path in the way that Gärdenfors describes [13, p. 221]:

> The prime problem is that the information received by the receptors is too rich and unstructured. What is needed is some way of transforming and organizing the input into a mode that can be handled on the conceptual or [according to Gärdenfors, more abstract] symbolic level. This basically involves finding a more *economic* form of representation: going from the subconceptual to the conceptual level usually involves a *reduction of the number of dimensions* that are represented.

Such salience for the conceptual agent must, seemingly, be grounded in something far more basic: most likely, something that applies to all cognitive agents. Here, *sensorimotor++* does not have much to offer beyond a promissory note and an appeal to the attempt by many in the enactive community to ground 'minimal' salience in the survival of the organism (see e.g. [52]): what is salient is what enables the organism to survive. As the agent becomes more complex and develops a somatosensory system, that system then comes to play a key role.

As I suggested in Section 3.2, talk of 'inside' and 'outside' with respect to cognitive phenomena are unhelpful at best, misleading at worst; so it is with the familiar distinction between the five 'external' senses of taste, touch, smell, sight, and hearing, and the 'internal' senses of interoception (awareness of bodily states) and proprioception (awareness of bodily positions and movements).[16] Certainly, on the perspective where the observer is pushed into the background (if not ignored altogether), it makes sense to talk of how the different sensory modalities are integrated in the brain; but on the competing perspective, where the observer is in the foreground and one's focus is on mind more than world, then it makes much more sense to talk about how one starts from undifferentiated experience, which then gets divvied out into the various modalities and – for the self-reflective conceptual agent – subsequently conceptualized into realms of 'internal' self and 'external' non-self/other/world.[17]

4.2 From Protoconcepts to Concepts

As noted in the previous section (see also Footnote 4), on Gallese and Lakoff's account [12], all concepts – even the most abstract ones – are nothing more than specific sensorimotor engagements, albeit with parts of those engagements suppressed; no additional cognitive mechanism is required. The example they make the most reference to is the concept of *grasp*. All that the concept amounts

[16] This distinction is clearer in some disciplines than others; as one of the anonymous reviewers of this paper noted, in biology, proprioception and touch are often seen to be closely related.

[17] I owe this point to psychologist Marek McGann, whom I have heard to make it in several conference presentations.

to, on any given usage, is a specific occasion of (physical) grasping, even if one does not physically carry it out: that is, the concept *is no more nor anything other than a variation on the familiar physical action.*

As I have suggested, such an approach seems inadequate for cognition in general; but, when it comes to conceptual cognition, its shortcomings are particularly striking. Here, one might well accuse Gallese and Lakoff of having chosen a deliberately 'concrete' concept; yet, like the gaping differences between so-called mental imagery and pictures such as those hanging on the wall, the concept of grasping – which allows grasping ideas and intentions as much as door handles or hammers – is, in many ways, not like an 'actual' physical grasping at all. In particular: to the extent that one's concept of grasping is a representation of grasping – and I have argued that, for the conceptually reflective agent, all concepts take on this representational aspect – the representation may have as much, or as little, to do with the represented as e.g. a representation of a dog, such as a painting of a dog, has to do with an actual dog. What goes for grasping goes for all the more abstract concepts like procrastination or *ennui* or the concept of concept itself. In keeping with the sensorimotor account, *sensorimotor++* and the *unified conceptual space theory* agree that, certainly, sensorimotor engagement is necessary to the foundation of even the most abstract of concepts; at the same time, it is not sufficient. *Pace* Gallese and Lakoff, some additional cognitive mechanism is required.

As the quote from Gärdenfors suggests, one of the key roles of concepts is to simplify: precisely to distance oneself from the world in order to better understand it. Concepts abstract away from the moment, from the particulars of context, from – *pace* Gallese and Lakoff and Berkeley – any particular application of them. They allow the agent to step back from the immediacy of their sensorimotor engagement; in which case, some additional mechanism or mechanisms is required, to actualize the path from meaninglessness to salience.

I propose that the unified conceptual space theory provides this 'representationally disposed' mechanism, described through its algorithm or 'recipe' for constructing and de-constructing concepts. Like Gallese and Lakoff's account, it rejects a prior ontological class/instance distinction; but, whereas Gallese and Lakoff (and Berkeley) reject classes in favour of instances, it ultimately rejects instances in favour of classes: to wit, any concept taken as a specific instance of something can contrastingly be understood as a class of yet more specific instances: so e.g. *dog* is a particular animal and, at the same time, a class of various breeds; while *my dog Fella* is a particular dog and, at the same time, a class of all my experiences and interactions with Fella.

One need not go so far as Fodor's *radical nativism* [9,11,23], by which most concepts[18] are innate[19], to allow that *something* innate is needed to kickstart the endless cycle of concept acquisition and application (see Section 3.1). What the *unified conceptual space theory* offers – rather more modestly than radical

[18] Fodor's target is *lexical concepts*, which, for Fodor, means most concepts.

[19] What precisely Fodor means by 'innate' is a subject of some controversy; for a good account of what he *probably* means, see [43, p. 230].

nativism! – is a small set of innate protoconcepts (or, if you will, protoconceptual abilities), along with the conjunctive and disjunctive connectors needed to bind them together or progressively partition them into sub-concepts.[20] These protoconcepts are suggested to consist of *proto-object, proto-property,* and *proto-action/event,* corresponding roughly to the grammatical categories, in English, of noun, adjective/adverb, and verb. They are not true concepts because they fail to meet all the usual desiderata of concepts as listed in the introduction: in particular, they are too few in number to be, of themselves, productive; and, being innate and therefore passively given to the agent, they are not under the agent's *endogenous control.* By one means or another[21], we seem predisposed to encounter the world and structure our understandings of it in terms of (concrete and abstract) objects, (concrete and abstract) happenings, and the properties of both – to the extent that it seems impossible to imagine encountering the world any other way. Given the appropriate environment and the appropriate sensorimotor engagements with that environment, these protoconcepts and connectors can, or so the *unified conceptual space theory* claims, give rise to the most richly structured of conceptual frameworks. In the language of *conceptual spaces theory,* they do so through the progressive partitioning of an initially minimally partitioned protoconceptual space (see the left-most illustration in Figure 1) – by extracting from perception patterns, patterns of patterns, and patterns of patterns of patterns, whilst disregarding or discarding others.

5 Conclusions

Like the original 2001 version and many of the subsequent formulations of sensorimotor theory, and in keeping with the enactive tradition, *sensorimotor++* attempts to resolve the seeming explanatory gap between mind and body, subjective experience and objective world; but it does so in a different way from the 2001 paper. What it comes down to, according to *sensorimotor++* and the *unified conceptual space theory* of which it is a part, is two different perspectives we move constantly – and for the most part non-reflectively – between. In the one, the observer is front and centre; in the other, the observer is pushed into the background or, seemingly, eliminated altogether. As the discussion has meant to imply, we cannot resolve the two perspectives into one, unified perspective because of our position within the explanatory loop; resolving the tension between perspectives would require stepping outside the loop.

Sensorimotor engagement – as it is commonly understood, in terms of the so-called 'external' senses – is necessary but not sufficient for understanding either

[20] For more on protoconcepts, as described within the *unified conceptual space theory,* see [39,36,40]. A yet more detailed account of protoconcepts is intended for an upcoming paper.

[21] I can, and do, choose to remain agnostic about what precisely is meant by 'innate': whether these protoconcepts are directly or indirectly specified in the genes or develop in the womb or something else again.

cognition in general or (my interest) conceptual cognition in particular. Incorporating interoception, proprioception, and emotion helps resist the extremes of externalism without falling into the trap of internalism. It strikes what I believe to be the ideal – if highly tensioned – balance between the two whilst suggesting how an account of salience might be grounded. Adding what I call a representationally disposed cognitive mechanism – described, in the *unified conceptual space theory,* by an algorithm for concept formation and evolution – actualizes that account. Representation – and therefore, *pace* Hutto, mental content – is a necessary part of any account of concepts and conceptual cognition that in any way tries to be complete, if only because, whenever we stop to reflect on our concepts *as* concepts – representations are what we find.

References

1. Berkeley, G.: Principles of Human Knowledge and Three Dialogues. Oxford University Press (1999), a Treatise Concerning the Principles of Human Knowledge was first published in 1710.
2. Brooks, R.A.: Elephants don't play chess. Robotics and Autonomous Systems 6, 3–15 (1990)
3. Brooks, R.A.: Intelligence without representation. Artificial Intelligence 47, 139–159 (1991)
4. Buhrmann, T., Paolo, E.D., Barandiaran, X.: A dynamic systems account of sensorimotor contingencies. Frontiers in Psychology 4, 1–19 (2013), doi:10.3389/fpsyg.2013.00285.
5. Chalmers, D.J.: Facing up to the hard problem of consciousness. In: Hameroff, S.R., Kaszniak, A.W., Scott, A. (eds.) Toward a Science of Consciousness: The First Tucson Discussions and Debates, pp. 5–28. MIT Press (1996)
6. Clark, A.: Rick: Towards a cognitive robotics. Adaptive Behavior 7(1), 5–16 (1999), `http://hdl.handle.net/1842/1297`
7. Damasio, A.: The Feeling of What Happens: Body, Emotion and the Making of Consciousness. Vintage (2000)
8. Evans, G.: Varieties of Reference. Clarendon Press (1982), edited by John McDowell
9. Fodor, J.A.: The Language of Thought. Crowell (1975)
10. Fodor, J.A.: Concepts: Where Cognitive Science Went Wrong. Clarendon Press, Oxford (1998)
11. Fodor, J.A.: LOT 2: The Language of Thought Revisited. Oxford University Press (2008)
12. Gallese, V., Lakoff, G.: The brain's concepts: The role of the sensory-motor system in conceptual knowledge. Cognitive Neuropsychology 22(3-4), 455–479 (2005)
13. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought. Bradford Books (2004)
14. Gibson, J.J.: The Ecological Approach to Visual Perception. Lawrence Erlbaum Associates (1986)
15. Grush, R., Churchland, P.: Gaps in Penrose's toilings. In: Metzinger, T. (ed.) Conscious Experience, pp. 185–214. Imprint Academic (1995)
16. Grush, R.: The emulation theory of representation: Motor control, imagery, and perception. Behavioral and Brain Sciences 27, 377–442 (2004), `http://escholarship.org/uc/item/15t2595z`

17. Harvey, I.: Untimed and misrepresented: Connectionism and the computer metaphor (CSRP 245), university of Sussex (UK) Cognitive Science Research Papers (CSRP) series (1992), http://www.sussex.ac.uk/Users/inmanh/harvey92untimed.pdf
18. Haugeland, J.: Artificial Intelligence: The Very Idea. MIT Press (1989)
19. Held, R., Hein, A.: Movement-produced stimulation in the development of visually guided behavior. Journal of Comparative and Physiological Psychology 56(5), 872–876 (1963)
20. Hofstadter, D.: Gödel, Escher, Bach: An Eternal Golden Braid. Penguin (2000), 20th Anniversary edition
21. Hutto, D.D.: Knowing what? radical versus conservative enactivism. Phenomenology and the Cognitive Sciences 4, 389–405 (2005)
22. Jaegher, H.D., Paolo, E.D., Gallagher, S.: Can social interaction constitute social cognition? Trends in Cognitive Science (2010) (in press)
23. Laurence, S., Margolis, E.: Radical concept nativism. Cognition 86, 25–55 (2002)
24. Machery, E.: Doing Without Concepts. Oxford University Press (2009)
25. Maturana, H.: Cognition. In: Hejl, P.M., Köck, W.K., Roth, G. (eds.) Wahrnehmung und Kommunikation, pp. 29–49. Peter Lang, Frankfurt (1978), http://www.enolagaia.com/M78bCog.html (with the original page numbering retained)
26. Maturana, H., Varela, F.J.: Autopoiesis and Cognition: The Realization of the Living (Boston Studies in the Philosophy of Science). Kluwer Academic Publishers (1980)
27. Maturana, H.R., Varela, F.J.: The Tree of Knowledge: The Biological Roots of Human Understanding. Shambhala, London (1992)
28. Millikan, R.: A common structure for concepts of individuals, stuffs, and real kinds: More mama, more milk, and more mouse. Behavioral and Brain Sciences 21, 55–100 (1998)
29. Morse, A., Ziemke, T.: The somatic sensory hypothesis (2010) (unpublished manuscript)
30. Newell, A.: Physical symbol systems. Cognitive Science 4(2), 135–183 (1980)
31. Noë, A.: Action in Perception. MIT Press (2004)
32. Noë, A.: Vision without representation. In: Nivedita, G., Madary, M., Spicer, F. (eds.) Perception, Action, and Consciousness: Sensorimotor Dynamics and Two Visual Systems, pp. 245–256. Oxford University Press (2010)
33. Novak, J., Canas, A.: The theory underlying concept maps and how to construct them. technical report, Florida Institute for Human and Machine Cognition (January 2008), http://cmap.ihmc.us/Publications/ResearchPapers/TheoryUnderlyingConceptMaps.pdf
34. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Beha 24, 939–1031 (2001)
35. O'Regan, K., Myin, E., Noë, A.: Sensory consciousness explained (better) in terms of 'corporality' and 'alerting capacity'. Phenomenology and the Cognitive Sciences 4, 369–387 (2005), doi:10.1007/s11097-005-9000-0.
36. Parthemore, J.: Concepts Enacted: Confronting the Obstacles and Paradoxes Inherent in Pursuing a Scientific Understanding of the Building Blocks of Human Thought. Ph.D. thesis, University of Sussex, Falmer, Brighton, UK (March 2011), http://sro.sussex.ac.uk/6954/1/Parthemore
37. Parthemore, J.: Of boundaries and metaphysical starting points: Why the extended mind cannot be so lightly dismissed. Teorema 30(2), 79–94 (2011)

38. Parthemore, J.: Representations, symbols, icons, concepts.. and why there are no mental representations. In: Proceedings of the Seventh Conference of the Nordic Association for Semiotic Studies, May 6-8 2011 (2013) (forthcoming)
39. Parthemore, J.: The unified conceptual space theory: An enactive theory of concepts. Adaptive Behavior (2013) (in press)
40. Parthemore, J., Morse, A.F.: Representations reclaimed: Accounting for the co-emergence of concepts and experience. Pragmatics & Cognition 18(2), 273–312 (2010)
41. Penrose, R.: Shadows of the Mind: A Search for the Missing Science of Consciousness. Oxford University Press (1994)
42. Perry, J.: Thought without representation. In: Proceedings of the Aristotelian Society, vol. 60, pp. 137–151 (1986)
43. Prinz, J.: Furnishing the Mind: Concepts and Their Perceptual Basis. MIT Press (2004)
44. Rosch, E.: Family resemblances: Studies in the internal structure of categories. Cognitive Psychology 7, 573–605 (1975)
45. Rosch, E.: Principles of categorization. In: Margolis, E., Laurence, S. (eds.) Concepts: Core Readings, ch. 8, pp. 189–206. MIT Press (1999)
46. Ryle, G.: The Concept of Mind. Penguin (1949)
47. Sharples, M.: How We Write: An Account of Writing as Creative Design. Routledge (1999)
48. Sonesson, G.: The mind in the picture and the picture in the mind: A phenomenological approach to cognitive semiotics. Lexia: Rivista di Semiotica 7-8, 167–182 (2011)
49. Steiner, P., Stewart, J.: From autonomy to heteronomy (and back): The enaction of social life. Phenomenology and the Cognitive Sciences 8, 527–550 (2009)
50. Stewart, J.: Cognition = life: Implications for higher-level cognition. Behavioural Processes 35(1-3), 311–326 (1995)
51. Thompson, E.: Mind in Life: Biology, Phenomenology and the Sciences of Mind. Harvard University Press (2007)
52. Thompson, E., Stapleton, M.: Making sense of sense-making: Reflections on enactive and extended mind theories. Topoi. 28(1), 23–30 (2009)
53. Varela, F.J.: Laying down a path in walking. In: Thompson, W. (ed.) Gaia: A Way of Knowing. Political Implications of the New Biology, pp. 48–64. Lindisfarne Press, Hudson (1987)
54. Varela, F.J., Thompson, E., Rosch, E.: The Embodied Mind: Cognitive Science and Human Experience. MIT Press (1991)

Conscious Sensation, Conscious Perception and Sensorimotor Theories of Consciousness

David Gamez

Department of Informatics, University of Sussex, Brighton BN1 9QJ
david@davidgamez.eu

Abstract. This article explores the hypothesis that the differences between our conscious sensations (color, sound, smell, etc.) could be linked to the different ways in which our senses process and structure information. It is also proposed that the organization of our conscious sensations into a conscious perception of a three-dimensional world could be linked to our mastery of sensorimotor contingencies. These hypotheses are supported by a number of observations, including the appearance of conscious sensations without motor action and the apparent failure of sensory substitution systems to generate visual sensations in congenitally blind subjects. The article discusses how the correlates of conscious sensation and perception could develop in the brain and some suggestions are put forward about how this account could be experimentally tested.

Keywords: consciousness, sensation, sensory substitution, perception, sensorimotor contingencies, correlates of consciousness.

1 Introduction

This article explores the hypothesis that the differences between our conscious sensations (color, sound, smell, etc.) could be linked to the different ways in which our senses process and structure information.[1] It also proposed that our sensorimotor interactions with our environment enable us to organize these sensations into a consciously perceived three-dimensional world. I will start with some definitions that will enable me to state these hypotheses more precisely:

- *Sensory contingencies*. Different senses have highly characteristic ways of processing information from the world. For example, the photoreceptors in the retina have a variety of response characteristics and spatial distributions, and they are wired up in complex ways to bipolar, horizontal, and ganglion cells, which work together to produce complex patterns of spikes in response to light. The other senses (sound, taste, smell, proprioception, etc.) also have unique spatiotemporal

[1] In this article I am using 'information' in a loose non-technical sense. In other work [1] I have suggested how Floridi's [2] approach could help us to develop better ways of identifying information in the brain.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_11,

response characteristics that process information in complex ways. The different ways in which the senses process incoming information will be referred to as *sensory contingencies*.

- *Sensorimotor contingencies*. Perception is an active process in which we move our eyes and body to acquire sensory information. The specific ways in which these sensorimotor patterns are structured are called *sensorimotor contingencies*. For example, when I am looking at a line the sensory information remains unchanged as I move my eyes along the line, and it alters when I move my eyes across the line.[2] When I move my hand from left to right, the visual image of my hand changes in a predictable way.
- *Conscious sensations*. I will use the term *conscious sensations* to describe the qualitative aspects of our experience, such as red, pain, birdsong, the taste of chocolate, and so on. Although 'qualia' is more commonly used to describe these phenomena, it is a controversial philosophical term, and 'conscious sensations' has the advantage that it is explicitly connected with the senses and can be more naturally contrasted with unconscious sensations.
- *Conscious perception*. This is our normal experience of perceiving our environment. While conscious sensations can appear without being referred to a cause or object (see Section 2.1), they are typically integrated into a consciously perceived three-dimensional world – for example, I see a red bird on the ground, smell it and hear the rustle of its feathers. A crude way of picturing the distinction between conscious sensation and conscious perception is that conscious perception provides an integrated three-dimensional framework that is 'filled in' with conscious sensations.

This article will explore the following hypotheses:

H1. Conscious sensations are correlated with sensory contingencies. The differences between the sensory channels' information-processing can be used to explain and make predictions about the qualitative differences between conscious sensations - for example, the difference between color and sound.[3]

H2. Conscious perception is correlated with our mastery of sensorimotor contingencies. Our interactions with the world teach us the relationships between motor outputs and sensory inputs. This enables us to structure our conscious sensations in a consciously perceived three-dimensional space centered on the body.[4]

The context and level of intensity of incoming sensory information determine whether it is experienced as conscious sensation (a sudden noise, pain, buzzing on the skin) or integrated with proprioceptive and motor information to produce conscious perception.

---

[2] This example is from O'Regan and Noë [3].

[3] A version of this hypothesis is defended by Keeley [4], who gives a useful summary of previous philosophical work on the differences between the senses.

[4] This account of the perception of the external world is similar to Aleksander's notion of depiction [5].

While some authors have suggested that there is an identity between a mastery of sensorimotor contingencies and consciousness [3, 6, 7], I am only focusing here on the weaker and less contentious claim that there might be a correlation between sensory or sensorimotor contingencies and the contents of consciousness. Since I am only examining the correlates of different types of content, I will set aside questions about the correlates of the level of consciousness. Given that a brain is conscious, I want to know why a particular pattern of activity is correlated with a conscious sensation of red instead of green, or with a sound instead of a smell?[5]

Some sensorimotor theories tend towards the view that patterns of activity in the brain *and world* are correlated with conscious states [6]. This fails to explain situations in which the contents of consciousness are disconnected from the world, such as dreams, out of body experiences, memories, hallucinations, or when the brain is stimulated with electrodes or TMS. In this article I will focus on the more plausible hypothesis that neural activity patterns in the brain are correlated with conscious contents. The brain learns these patterns by interacting with the world; when they are reactivated in a trained brain they can potentially be correlated with conscious contents independently of the current environment (see Section 3).[6]

The first part of this article puts forward a number of observations that support a dissociation between conscious sensation and sensorimotor contingencies. I will then suggest how the correlates of conscious sensation and perception could emerge in a developing brain. Some ways of experimentally testing H1 and H2 are then put forward, followed by a discussion of how they relate to the sensorimotor theories of O'Regan and Noë. The article concludes with some implications of these hypotheses.

2 Conscious Sensations and Sensory Contingencies

This section highlights a number of situations in which conscious sensations occur without motor action, which suggests that they are unlikely to be correlated with or identical to sensorimotor contingencies. The research on sensory substitution systems suggests that conscious sensations are linked to sensory contingencies and sensorimotor contingencies are associated with our conscious perception of a three-dimensional world.

---

[5] This is broadly in agreement with Chalmers' definition of the neural correlates of the content of consciousness: "An NCC (for content) is a minimal neural representational system N such that representation of a content in N is sufficient, under conditions C, for representation of that content in consciousness." ([8], p. 31).

[6] See [3, 6] for attempts to resist this position, which in my view are unconvincing. While O'Regan [7] acknowledges that brain stimulation could induce conscious sensations (p.108), he claims that conscious sensations are qualities of ongoing interactions between the brain and world, not some kind of essence that is generated by the brain. Some of the commentary on [3] discusses this issue, which is covered in more detail in Section 5.1.

2.1 Conscious Sensations without Action

A number of observations support a dissociation between conscious sensations and motor action. First, we often experience conscious sensations without having made a motor action. For example, I am sitting passively in the dark and a bright light is suddenly switched on. I initially experience a raw sensation of blinding light, and then my eyes adapt, I start to saccade and I perceive a structured world. The phenomenology is similar with a sudden noise: at first I am startled by the noise, which swamps my senses and I am consumed by the sensation of it; some moments later I locate the source and nature of the sound.

Second, experiments have shown that conscious sensations can be produced by very brief stimuli – for example, a 1ms flash of light can cause a conscious visual sensation; a conscious auditory sensation can be elicited by a 1ms auditory click [9]. The duration of these stimuli is much less than the timescale of motor actions that could actively explore them (the eyes saccade every ~200ms; ear and head movements take much longer).

Third, many conscious body sensations, such as heartburn and pain, do not have an obvious motor component – they just happen to us and we do not have to move or do anything to make them happen or go away. All of these observations suggest that motor action is not necessary for conscious sensations.[7]

2.2 Direct Brain Stimulation Produces Conscious Sensations

Conscious sensations can be evoked by directly stimulating the brain of a passive subject. For example, a blow to the head, a TMS pulse or electrode stimulation produces phosphenes, memories, body sensations and sounds [10, 11].

2.3 Sensory Substitution Systems

The phenomenology of people's use of sensory substitution systems is easily explained by a distinction between conscious sensation and perception. For example, in a tactile visual substitution system (TVSS) a two-dimensional array of vibrators is placed on the body[8] and connected to visual information from a camera that is typically mounted on the head [12]. If the subject does not move, they typically experience a buzzing on the surface of their body. While little or no explanation of this buzzing sensation is offered by a sensorimotor theory of consciousness, it becomes easier to understand if there is a correlation between somatosensory contingencies and conscious buzzing sensations (H1).

---

[7] O'Regan [7] acknowledges that vision can occur without action, and claims that being poised for action is enough, even if the action is not executed (p. 67).

[8] In some TVSSs a two dimensional array of electrodes is placed on the tongue: a small voltage produces a tingling sensation.

When subjects are allowed to move the camera, the buzzing sensation is transformed into perception of a spatially organized world.[9] However, although congenitally blind subjects can use a TVSS to perceive a three-dimensional environment, this is not enough to give them a conscious sensation of light. Visually handicapped people expressed disappointment with these devices for this reason [13].[10]

The absence of conscious visual sensations in congenitally blind subjects is supported by a study in which normally sighted and congenitally blind subjects were trained on a tongue-mounted TVSS. Before training the blind subjects did not report any tactile sensations on their tongues when TMS was applied to their visual cortex. After training the TMS caused some of the blind subjects to experience somatopically organized tactile sensations on their tongues. The application of TMS to the same brain areas of the normally sighted subjects caused them to experience visual phosphenes both before and after the training [15]. This suggests that the TVSS training increased the blind subjects' ability to perceive the world through tactile sensations on their tongues, but it did not give them the conscious sensation of light.[11]

While H1 predicts that congenitally blind subjects using a TVSS will not experience conscious visual sensations, a limited amount of visual experience would be expected in non-congenitally blind or normal subjects. Once these subjects have learnt to use the TVSS, they can use their memory and imagination to generate visual images of the world that they are perceiving through the buzzing sensations. There are also a considerable number of cross-sensory connections in the brain, which could lead to visual areas being activated in response to other sensory stimulation - a phenomena that appears most strongly in synaesthesia.

The limitations of sensory substitution systems also support a dissociation between conscious sensations and sensorimotor contingencies. While some success has been achieved with devices that substitute vision using audio or somatosensory stimulation, no-one has created a taste or pain sensory substitution device, and it is very difficult to imagine how this could work. We might try to build a vision-taste substitution system by giving a person a tasteless object to chew while they look at a display showing the sensory patterns that are produced by the taste receptors [18]. However, it seems inconceivable that such a system could evoke the sensation of bitterness or

---

[9] Subjects have to spend some time practicing with the device before this occurs.

[10] The first person report of a congenitally blind person's experience with a TVSS [14] is sometimes cited as a counter-example to this claim. However, Guarniero's description of 'seeing' with the device is much more akin to spatial perception, than encountering what for him would be a novel conscious sensation: "As I have noted at the beginning of this paper, I have used the word 'see' for lack of a better. The difficulty is not merely one of vocabulary; rather it is a conceptual one. Very soon after I had learned how to scan, the sensations no longer felt as if they were located on my back, and I became less and less aware that vibrating pins were making contact with my skin. By this time objects had come to have a top and bottom; a right side and a left; but no depth – they existed in an ordered two-dimensional space, the precise location of which has not yet been determined." (p. 104)

[11] Similar interpretations of TVSS experiments are given by Keeley [4], Block [16] and Prinz [17].

sweetness in a subject, and similar problems exist with the substitution of pain or smell.[12,13]

2.4 Ockham's Razor

Some theories of consciousness claim that sensorimotor patterns can be used to explain the differences between our senses (for example, [3]). However, if sensory contingencies contain enough information to differentiate the senses, then it might not be necessary to include motor actions in our explanations as well. If we are looking for a *minimal* set of spatiotemporal structures in the brain that are correlated with conscious sensations, then we should start with sensory contingencies, and carry out experiments to see if these provide enough information to make accurate predictions about conscious contents (see Section 4).

2.5 Conscious Sensations without Sensorimotor Contingencies

In some situations people experience conscious sensations that they cannot systematically integrate with their actions [19]. For example, people who have had their sight restored after a cataract operation often have mixed-up sensations of color, movement and light and do not know how to use their eye movements to perceive the world [20]. A similar situation occurs when people wear inverting glasses that disrupt the normal relationship between changes in the visual world and eye and body movements. When subjects put on these glasses they are initially extremely disoriented and find it very difficult to perceive or interact with the world. Gradually they mastery the sensorimotor contingencies of the inverting glasses and learn to perceive and act in the world again [21].

While Noë [19] describes these situations as 'experiential blindness', it is more accurate to describe them as cases of perceptual blindness, since the subjects can experience a variety of color, light and movement sensations. They are perceptually blind because they do not know the sensorimotor rules that would enable them to coordinate their experiences into a consciously perceived world. People become experientially blind (lose conscious visual sensations) if their sensory apparatus (eyes, optic nerves, visual cortex, etc.) is damaged.

3 The Correlates of Conscious Sensation and Perception

The observations outlined so far suggest that there is a connection between conscious sensations and sensory contingencies in the brain. However, conscious sensations

[12] One way of explaining the limited success of sensory substitution systems is that they provide information that can be accessed through a number of different sensory channels, but they cannot substitute information that is only present in a single sensory channel.

[13] Some of these points were made by Fiona MacPherson in her talk 'Sensory Substitution and Augmentation: Introduction to the Issues' at the conference on Sensory Substitution and Augmentation, British Academy, London, 26th March 2013.

cannot be directly correlated with the low-level processing of the individual senses because they can occur without the activation of the sensory hardware – for example, conscious visual sensations can be triggered by TMS independently of the retina. A more plausible hypothesis is that the cortical areas connected to each sensory channel learn to respond to the sensory contingencies of that channel, which enables them to reproduce the sensory contingencies in the absence of stimulation from the senses. Research on the neural correlates of consciousness has indicated that low level sensory areas, such as V1, are only weakly correlated with consciousness [22]. This suggests that the brain areas which are correlated with conscious sensations are likely to be higher up the sensory processing hierarchies.

The developing cortex is only roughly wired up using chemical gradients and other mechanisms, and so the differences between our conscious sensations are unlikely to be linked to genetically wired connection patterns. It is more likely that the developing cortex learns to respond to incoming sensory patterns using synaptic pruning and other processes, which leads to substantially different structures and response characteristics in the different sensory areas [23]. In the womb the data that it is available for this learning process includes random noise in the sensors (retina, cochlea, etc.) and environmental stimulation (sound, light, taste and smell). Once a cortical area has learnt to respond to a set of sensory contingencies, artificial stimulation of the area will produce a noisy version of the corresponding sensation – for example, TMS stimulation of trained visual cortex leads to phosphenes.

The structuring of the cortex in response to sensory stimulation patterns has been experimentally demonstrated in ferrets and hamsters. For example, in work carried out by Sur et al. [23] the visual pathways in neonatal ferrets were redirected to the auditory cortex. Many of the neurons in the rewired auditory cortex developed visual response characteristics similar to those in V1, although there were some differences between orientation selective cells in V1 and the rewired cortex.[14] Similar results have been obtained in rewired hamsters [24], whose visually guided behavior was similar to controls [25].[15]

The development of the mechanisms linked to conscious perception probably occurs during late embryological development and after birth, when the child learns how its motor actions generate predictable patterns of sensory information. This probably strengthens the connections between the learnt sensory contingency patterns, proprioceptive areas and motor output areas.

---

[14] Some of the differences between rewired visual cortex and normal visual cortex could be due to the fact that the rewiring was carried out postnatally.

[15] It might be objected that people who have had their sight restored by a cataract or cornea operation have conscious visual sensations, although their cortex has apparently had no chance to learn to respond to visual information during early childhood. However, this type of operation is only carried out on patients with functional retinas [20]. This type of subject perceives a limited amount of light through the surface of the eye, which is enough to stimulate the cortex with visual sensory contingencies.

4 Testing H1 and H2

The first step in the testing of hypotheses H1 and H2 is the recording of cortical activity in which the sensory and sensorimotor contingencies that might be correlated with conscious contents can be found. A variety of techniques can be used to identify mathematical regularities in this data, which can be used to generate testable predictions about conscious sensations and perception.

4.1 Data

The first step in the recording of data is the identification of the parts of the brain that contain the sensory contingencies which could be correlated with conscious content. These can be identified by experiments on the correlates of consciousness, using binocular rivalry, the subliminal presentation of stimuli or other techniques [22]. Once appropriate areas of the brain have been identified we need to record the sensory and sensorimotor contingency patterns that the cortex has learnt as a result of its stimulation from the senses.

This type of data is difficult to record in humans because of the low spatial and/or temporal resolution of fMRI and EEG, and electrodes can only be implanted in a limited number of sites when patients are being operated on for other reasons. In animals optogenetic techniques are reaching the point at which they can record from up to 100,000 neurons at 1 Hz from zebrafish larvae [26], and it is becoming possible to record from a few tens of thousands of neurons close to the surface of the cortex of a mammalian brain. The problem with using data recorded from animals is that their sensory contingencies are likely to be very different from our own, and so predictions about conscious sensations based on this data are likely to be specific to the animals that it is taken from.

A more systematic way of understanding how the cortex learns sensory contingencies would be to prepare samples of embryological cortical tissue and expose them to patterns of activity from different senses. The stimulation patterns could be generated using spike conversion libraries, which have been developed for visual, proprioceptive and auditory data [27, 28]. If the picture sketched out in Section 3 is correct, these pieces of cortex should self-organize in response to the incoming data and exhibit different spontaneous activity patterns once they had been trained. To minimize the difficulties of working with in vitro tissue, this approach could be prototyped on simulated neural tissue – the models developed by Izhikevich and Edelman [29] or Markram [30] would be good starting points for this work.

4.2 Mathematical Regularities in the Data

Once the sensory and sensorimotor contingency patterns have been recorded we need to find a way of describing these patterns that is generalizable across individuals and can make accurate predictions about conscious sensations and perceptions. The research on brain reading using fMRI has used inferred models and statistics to make fairly accurate predictions about conscious contents, to the extent of decoding people's dreams [31] or reconstructing the videos they are watching [32]. However, these techniques are based on data with low spatial and temporal resolution, and they typically have to be fine tuned for each individual. This suggests that they have not completely captured the correlates of conscious contents in the brain.

An alternative way of tackling this problem would be to take inspiration from physics and look for mathematical regularities in the sensory and sensorimotor patterns that are not specific to any individual person. These could be expressed using sets of differential equations, category theory or some other mathematical formalism. The mathematical equations in physics are typically written down by an expert scientist, and this approach has been taken in consciousness science by Tononi [33], who has developed algorithms for generating a mathematical structure that is predicted to correspond to the contents of consciousness.

The central problem with the use of an expert scientist to compose the equations is that the regularities that develop in the cortex in response to sensory and sensorimotor contingencies might be too complicated to be captured in equations that are written down by a human. To avoid this problem, machine learning techniques could be used to infer the equations from recordings of cortical data. A computational approach to the discovery of scientific knowledge has shown promise in a number of areas [34-36], and it could be a good way of identifying potentially complex regularities in brain activity patterns that are correlated with conscious sensation and perception.

A more radical approach would be to develop a mathematical model of how data is transformed by the senses and learnt by the cortex [37]. While it would be more elegant to infer the sensory contingencies from the structure of the sensory apparatus, the ways in which the cortex learns to respond to this information might be too complex for an analytical mathematical treatment.

4.3 Predictions

Once the sensory and sensorimotor contingency patterns have been identified and mathematically described they could be used to make predictions about conscious sensations and perceptions in humans.[16] One approach would be to use the inferred mathematical regularities to make predictions about how sensory and sensorimotor contingencies would appear in fMRI or EEG data. Ideally the model should be able to specify how the sensory contingencies corresponding to different colors, sounds, etc would appear in an fMRI or EEG recording, which would enable testable predictions to be made about subjects' first person reports.

[16] Predictions about conscious content in animals are difficult to test. While they can generate a behavioral output in response to a stimulus that is assumed to be conscious, they have no way of describing their conscious contents.

5 H1 and H2 in Relation to Other Sensorimotor Theories

While the hypotheses of this paper have been influenced by O'Regan and Noë's work on sensorimotor theory [3, 6, 7, 19], there are important differences between H1 and H2 and their account of conscious experience. O'Regan and Noë would be likely to broadly agree on the following points:

1. The brain does not contain a rich set of internal representations of the world.
2. Conscious sensations are linked to sensorimotor contingencies.
3. Conscious perceptions are linked to sensorimotor contingencies.
4. While the brain plays a role in conscious sensation and perception, the world is likely to be necessary for some types of conscious experience.
5. Sensorimotor theory can explain the link between physical world and consciousness.

In this paper I have argued that conscious perception is linked to our mastery of sensorimotor contingencies (H2) and I have focused on a correlates-based approach to consciousness that does not seek to explain the relationship between the physical world and consciousness. I also think that the physiology of the eye and research on change blindness and inattentional blindness (for example, [38, 39]) demonstrate that we do not have a rich set of internal representations that accurately track the world. This leaves two main points of disagreement: the extent to which brain activity is correlated with consciousness and the link between sensorimotor contingencies and conscious sensation. I will consider these in turn.

5.1 Consciousness and the Brain

A significant difference between the position outlined in this paper and sensorimotor theorists, such as O'Regan and Noë, is the idea that conscious sensation and perception are correlated with brain activity alone, rather than with the brain and world (see Section 1). For example, in their joint paper O'Regan and Noë [3] claim:

> There can therefore be no one-to-one correspondence between visual experience and neural activations. Seeing is not constituted by activation of neural representations. Exactly the same neural state can underlie different experiences, just as the same body position can be part of different dances. (p. 966)

A less radical position can be found in Noë's [19] later work: "A reasonable bet, at this point, is that *some* experience, or some features of some experiences, are, as it were, exclusively neural in their causal basis, but that full-blown, mature human experience is not." (p. 218). To discuss the role of the brain in relation to O'Regan and Noë's theories, I will distinguish between two types of conscious experience:

1. *Online conscious perception*. Detailed information from the environment can be accessed on demand. Colors are vivid; sounds are loud and clear; objects are stable.

2. *Offline conscious experiences*. These include experiences induced by brain stimulation, dreams, imagination, phantom limbs, out-of-body experiences and hallucinations Offline conscious experiences are often unstable, low resolution and low intensity, and their content is independent of the state of the environment. Conscious sensations are present in some or all of the sensory modalities and the states are weakly perceptual – for example, we see objects but cannot interact with them in a systematic way.

Although there is some overlap between these two types of experience (for example, people can mistake phantom limbs for actual limbs [40]), they are reasonably distinct categories for most people. Since the first type of experience is typically produced as the brain interacts with its environment, it is difficult to prove that it is just correlated with brain activity. In the second type of experience, conscious sensations and a limited form of conscious perception occur without any interaction between the brain and its environment, which suggests that the second type of experience is correlated with brain activity alone.

While it has not been proved that online conscious perceptual experiences are solely correlated with brain activity, it seems reasonable to make this into a working assumption, which can be tested using the experimental approach outlined in Section 4. Noë's bet that this type of conscious experience is not exclusively correlated with neural activity is a different working hypothesis, which can also be experimentally tested.

5.2 Conscious Sensations and Sensorimotor Contingencies

A substantial amount of evidence was presented in Section 2 which strongly suggests that conscious sensations can occur without motor action. This supports a dissociation between conscious sensations and sensori*motor* contingencies. This would be falsified if conscious sensations could be produced without the activation of sensory contingency patterns in the cortex - for example, if the conscious sensation of red could be induced by a TVSS.

Section 2.3 set out evidence in favor of the claim that no device capable of inducing conscious sensations has been created. While O'Regan [7] cites Guarniero's experiences with the device as a possible counterexample, I argued in Footnote 10 that Guarniero's description of 'seeing' with the device is much more akin to spatial perception, than encountering what for him would be a novel conscious sensation. O'Regan [7] also suggests that the limited capabilities of the current sensory substitution systems are likely to be preventing them from inducing conscious sensations:

> …current techniques of visual-to-tactile and visual-to-auditory substitution are a long way from the goal of achieving a real sense of vision. Using tactile or auditory stimulation, it is possible only to provide a few aspects of normal visual impressions, like the quality of being "out there" in the world, and of conveying information about spatial layout and object form. But the "image-like" quality of vision still seems far away. Indeed, because the eye is such a

> high-resolution sensor, it will probably never be possible to attain true substitution of the image-like quality of vision. (pp. 142-3)

O'Regan goes on to discuss other substitution systems, such as a glove designed to provide touch sensations for leprosy patients [41] and vestibular substitution systems. In the glove experiments, subjects experienced touch sensations on the forehead as if they were on the glove However, this is not an example of sensory *substitution* because it only altered the location of a conscious sensation - it did not provide the information from one sense through a different sensory channel. The apparent success of vestibular substitution systems also does not count against H1 because vestibular information is not obviously correlated with conscious sensations.

It might be claimed that a device capable of sensory substitution could be created, and that the possibility of this device invalidates H1. However, people have very different intuitions in this area and imagination and thought experiments are rarely an accurate guide to what is possible in the real world. While I agree that H1 would be invalidated by an actual sensory substitution device that induces conscious sensations, the question cannot be decided by people's competing intuitions about whether this could or could not be constructed.

6 Discussion and Conclusions

This article has argued that the qualitative aspects of conscious experience (conscious sensations) are correlated with sensory contingencies (H1), and that the organization of conscious sensations into a three-dimensional consciously perceived world is linked to our mastery of sensorimotor contingencies (H2). I have suggested that the sensory and sensorimotor patterns that are correlated with consciousness are located in the cortex, which learns to respond to the output of the different sensory channels. The hypotheses of this paper can be tested by using machine learning techniques to infer mathematical descriptions of sensory and sensorimotor contingency patterns, which can be used to generate predictions about conscious sensations and perceptions.

This article's approach to conscious content could extend and complement previous theories about neural representation that have been put forward. For example, it has been suggested that the place of a neuron in a hierarchy or population determines its representational content, and there are a number of interpretations of the neural code (for example, firing rate or spike timing relationships). The problem with these theories of neural representation is that they are generally based on correlations between neural activity and states of the world, which limits their ability to systematically account for differences in conscious contents – for example, why pattern P in population X is correlated with a red sensation and pattern Q in population Y is correlated with the sound of a police siren. In this article I have put forward a more theory-driven approach, which could improve our ability to identify correlations between neural activity and conscious sensations and perceptions.

The distinction between sensory and sensorimotor contingencies can help us to understand some of the limitations of sensory substitution devices. While they can successfully encode sensorimotor contingencies, they cannot alter or replace the contingent ways in which the sensory channels process information, and so people using them do not have the conscious sensations that correspond to the substituted sense (vibrations on the tongue do not induce conscious visual sensations). To effectively substitute conscious sensations, it would be necessary to model the sensory contingencies of a particular sense and feed the information directly into the cortex.

While sensory substitution systems have succeeded in transforming and augmenting our current sensations and perceptions, it is debatable whether they have been able to generate novel sensations.[17] According to the hypotheses put forward in this article, a novel conscious sensation could be induced in a subject by a sensory channel with unique sensory contingencies. This would have to be directly connected to the cortex, which would have to learn to respond to it over an extended period. Although optogenetics or implanted electrodes would be ideal for this task, there are few situations in which these can be used in human subjects. A more practical method would be to use focused ultrasound [43] or high-definition transcranial direct current stimulation to deliver an appropriate signal to the brain.[18] Alternatively, it might be possible to use an inverse model to cancel out the sensory contingencies of a particular channel. For example, perhaps one could develop an inverse model of the retina and early visual cortex and apply this to visual data to cancel out the effects of visual processing. This might enable data free of visual sensory contingencies to be passed directly to the brain. Further research is needed to determine the feasibility of this approach.

This article has focused on the possibility that conscious sensations could be correlated with sensory contingencies. I have not attempted to address the ‘hard’ problem of consciousness or tried to *explain* the relationship between patterns of brain activity and conscious sensations. Elsewhere I have argued that our inability to imagine the physical world will make it impossible to develop an intuitively satisfying explanation of the relationship between consciousness and brain activity.[19] However, empirical work on the correlates of consciousness could lead to the development of mathematically formulated theories that can map with high accuracy between brain activity and conscious states. This might make us more willing to abandon our desire for intuitively satisfying explanations in consciousness science, just as we have given up hope of intuitively satisfying explanations in quantum mechanics.

---

[17] It has been suggested that the the feelSpace belt could provide a qualitatively new perceptual experience [42]. However, the first person reports suggest that subjects’ existing sense of space, location and landmarks was expanded and made more accurate, not that an entirely novel conscious sensation was created.

[18] This technology is being developed by Soterix: `http://soterixmedical.com/hd-tdcs`. The spatial resolution would probably have to be improved before it could be used for sensory substitution.

[19] David Gamez, ‘The Hard and the Real Mind Body Problem’, unpublished. Available at: `http://www.davidgamez.eu/papers/Gamez_MindBodyProblem.pdf`

Further research is also needed to determine whether *all* sensory contingency patterns are correlated with conscious sensations[20] and why some senses are richer and more vivid than others.[21] Finally, the hypotheses put forward in this article could help us to develop a better understanding of the conscious contents of infants, animals and artificial systems, whose sensory and sensorimotor contingencies are very different from our own. It could also help us to address some of the limitations of the current work on the neural correlates of consciousness that were identified by Noë and Thompson [45].

Acknowledgements. This work was supported by Barry Cooper's grant from the John Templeton Foundation (ID 15619: 'Mind, Mechanism and Mathematics: Turing Centenary Research Project'). I would also like to thank Anil Seth and the Sackler Centre for Consciousness Science at the University of Sussex for hosting me as a Research Fellow during this project.

References

1. Gamez, D.: Information and Consciousness. Etica & Politica / Ethics & Politics XIII, 215–234 (2011)
2. Floridi, L.: Philosophical Conceptions of Information. In: Sommaruga, G. (ed.) Formal Theories of Information. LNCS, vol. 5363, pp. 13–53. Springer, Heidelberg (2009)
3. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, 939–973 (2001)
4. Keeley, B.L.: Making Sense of the Senses: Individuating Modalities in Humans and Other Animals. The Journal of Philosophy 99, 5–28 (2002)
5. Aleksander, I.: The world in my mind, my mind in the world. Imprint Academic, Exeter (2005)
6. Noë, A.: Out of our heads. Hill and Wang, New York (2009)
7. O'Regan, J.K.: Why red doesn't sound like a bell: Understanding the feel of consciousness. Oxford University Press, New York (2011)
8. Chalmers, D.: What Is a Neural Correlate of Consciousness? In: Metzinger, T. (ed.) Neural Correlates of Consciousness, pp. 17–39. MIT Press, Cambridge (2000)
9. Pockett, S.: How long is "now"? Phenomenology and the specious present. Phenomenol Cogn. Sci. 2, 55–68 (2003)
10. Brindley, G.S., Lewin, W.S.: The sensations produced by electrical stimulation of the visual cortex. J. Physiol. 196, 479–493 (1968)
11. Penfield, W.: Some mechanisms of consciousness discovered during electrical stimulation of the brain. Proceedings of the National Academy of Sciences 44, 51–66 (1958)
12. Bach-y-Rita, P.: Brain mechanisms in sensory substitution. Academic Press, New York (1972)

[20] Humans might have a vomeronasal organ whose sensory contingencies are not correlated with conscious sensations [4, 44].

[21] The resolution of the different senses would be one place to look for an explanation.

13. Lenay, C., Gapenne, O., Hanneton, S., Marque, C., Genouëlle, C.: Sensory Substitution: Limits and Perspectives. In: Hatwell, Y., Streri, A., Gentaz, E. (eds.) Touching for Knowing. John Benjamins Publishing, Amsterdam (2003)
14. Guarniero, G.: Experience of tactile vision. Perception 3, 101–104 (1974)
15. Kupers, R., Fumal, A., de Noordhout, A.M., Gjedde, A., Schoenen, J., Ptito, M.: Transcranial magnetic stimulation of the visual cortex induces somatotopically organized qualia in blind subjects. Proc. Natl. Acad. Sci. U S A 103, 13256–13260 (2006)
16. Block, N.: Tactile sensation via spatial perception. Trends Cogn. Sci. 7, 285–286 (2003)
17. Prinz, J.J.: Putting the Brakes on Enactive Perception. Psyche 12 (2006)
18. Chandrashekar, J., Hoon, M.A., Ryba, N.J., Zuker, C.S.: The receptors and cells for mammalian taste. Nature 444, 288–294 (2006)
19. Noe, A.: Action in perception. MIT, Cambridge (2004)
20. Gregory, R.L., Wallace, J.G.: Recovery from Early Blindness - A Case Study. Heffers, Cambridge (1963)
21. Kohler, I., Fiss, H.: The formation and transformation of the perceptual world. Translated by Harry Fiss. Introduction by James J. Gibson. International Universities Press, New York (1964)
22. Tononi, G., Koch, C.: The neural correlates of consciousness: an update. Ann. N. Y. Acad. Sci. 1124, 239–261 (2008)
23. Sur, M., Angelucci, A., Sharma, J.: Rewiring cortex: the role of patterned activity in development and plasticity of neocortical circuits. J. Neurobiol. 41, 33–43 (1999)
24. Ptito, M., Giguère, J.-F., Boire, D., Frost, D.O., Casanova, C.: When the auditory cortex turns visual. In: Casanova, C., Ptito, M. (eds.) Vision: From Neurons to Cognition. Elsevier Science, Amsterdam (2001)
25. Frost, D.O., Boire, D., Gingras, G., Ptito, M.: Surgically created neural pathways mediate visual pattern discrimination. Proc. Natl. Acad. Sci. U S A 97, 11068–11073 (2000)
26. Ahrens, M.B., Keller, P.J.: Whole-brain functional imaging at cellular resolution using light-sheet microscopy. Nature Methods (2013)
27. Gamez, D., Fidjeland, A.K., Lazdins, E.: iSpike: A Spiking Neural Interface for the iCub Robot. Bioinspiration and Biomimetics 7, 025008 (2012)
28. Fontaine, B., Goodman, D.F., Benichoux, V., Brette, R.: Brian hears: online auditory processing using vectorization over channels. Front Neuroinform 5, 9 (2011)
29. Izhikevich, E.M., Edelman, G.M.: Large-scale model of mammalian thalamocortical systems. Proc. Natl. Acad. Sci. U S A 105, 3593–3598 (2008)
30. Markram, H.: The blue brain project. Nat. Rev. Neurosci. 7, 153–160 (2006)
31. Horikawa, T., Tamaki, M., Miyawaki, Y., Kamitani, Y.: Neural Decoding of Visual Imagery During Sleep. Science (2013)
32. Nishimoto, S., Vu, A.T., Naselaris, T., Benjamini, Y., Yu, B., Gallant, J.L.: Reconstructing visual experiences from brain activity evoked by natural movies. Curr. Biol. 21, 1641–1646 (2011)
33. Tononi, G.: Consciousness as integrated information: a provisional manifesto. Biol. Bull. 215, 216–242 (2008)
34. Sparkes, A., Aubrey, W., Byrne, E., Clare, A., Khan, M.N., Liakata, M., Markham, M., Rowland, J., Soldatova, L.N., Whelan, K.E., Young, M., King, R.D.: Towards Robot Scientists for autonomous scientific discovery. Automated Experimentation 2, 1 (2010)
35. Džeroski, S., Todorovski, L. (eds.): Computational Discovery 2007. LNCS (LNAI), vol. 4660. Springer, Heidelberg (2007)
36. Schmidt, M., Lipson, H.: Distilling free-form natural laws from experimental data. Science 324, 81–85 (2009)

37. Gamez, D.: From Baconian to Popperian Neuroscience. Neural Systems and Circuits 2, 2 (2012)
38. Rensink, R.A., O'Regan, J.K., Clark, J.J.: To See or Not to See: The Need for Attention to Perceive Changes in Scenes. Psychological Science 8, 368–373 (1997)
39. Simons, D.J., Chabris, C.F.: Gorillas in our midst: sustained inattentional blindness for dynamic events. Perception 28, 1059–1074 (1999)
40. Melzack, R.: Phantom limbs. Scientific American 266, 120–126 (1992)
41. Bach-y-Rita, P., Kercel, S.W.: Sensory substitution and the human-machine interface. Trends Cogn. Sci. 7, 541–546 (2003)
42. Nagel, S.K., Carl, C., Kringe, T., Martin, R., Konig, P.: Beyond sensory substitution–learning the sixth sense. J. Neural. Eng. 2, R13–R26 (2005)
43. Bystritsky, A., Korb, A.S., Douglas, P.K., Cohen, M.S., Melega, W.P., Mulgaonkar, A.P., DeSalles, A., Min, B.K., Yoo, S.S.: A review of low-intensity focused ultrasound pulsation. Brain Stimul. 4, 125–136 (2011)
44. Meredith, M.: Human vomeronasal organ function: a critical review of best and worst cases. Chem. Senses. 26, 433–445 (2001)
45. Noë, A., Thompson, E.: Are There Neural Correlates of Consciousness? Journal of Consciousness Studies 11, 3–28 (2004)

Basic Pretending as Sensorimotor Engagement?

Lessons from Sensorimotor Theory for the Debate on Pretence

Zuzanna Rucinska

School of Humanities, University of Hertfordshire, Hatfield, United Kingdom
z.rucinska@hotmail.com, z.a.rucinska@herts.ac.uk

Abstract. This paper explores whether the sensorimotor theory of perception (SMTP) might contribute to a de-intellectualized understanding of pretence. It applies SMTP to Currie's [3], [4] notion of perceptual *seeing-in* that underlies the capacity to make imaginative transformations (*seeing-as*). This account bypasses manipulation of representational contents *off-line*, and argues that the relevant work might done by *on-line*, *sensory imaginings* stemming directly from perception. This novel position is supported with augmented theory of affordances and an account of directly perceived meaning. Ultimately, the paper proposes a less intellectualist approach than Currie's to object-substitution pretend play of young children, setting the stage for an enactive theory of basic pretence.

Keywords: Pretence, Pretending, Sensorimotor, Enactivism, Affordances, Imagination.

1 Introduction

From at least 18 months of age, when their use of language is still primitive, children engage in spontaneous pretence, as fun (with no ulterior motive, like an intention to deceive): they pretend that a banana is a telephone and that they are talking to it, that there is a dragon under the bed when there is nothing, that a doll's face is dirty when it is clean [11]. This behavior seems likely to be the primitive precursor to highly sophisticated fictions that the word 'fiction' naturally evokes: literary fictions are most salient, followed by plays and movies, arguably followed by painting and sculpture [20, 1].

This paper will take the case of banana-phone object-substitution play of pre-verbal children as genuine example of pretence and a paradigm case of basic pretend activity that can extend to more complex types of play. Pretend play, traditionally defined as *symbolic play*, has been taken by mainstream theories to require representational capacities of some kind.[1] At one end of the spectrum is the most

[1] Some even consider full-blown linguistic capacities to be necessary [9]. In this paper I will assume that the linguistic and conceptual capacities of 18 month olds are too limited for manipulation of propositions like 'this banana is now a phone', while at the same time that 18 month olds do engage in genuine pretence, thereby excluding for the sake of argument the possibility of a linguistic account of pretence.

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*, Studies in Applied Philosophy, Epistemology and Rational Ethics 15,

DOI: 10.1007/978-3-319-05107-9_12,

hyper-intellectualist, metarepresentational theory of pretence [11], which requires the skill to think of the banana as a phone in order to engage in the banana-phone pretence. At the other end, there are less-intellectualist accounts that appeal to imaginative or simulative activity that minimally only require pretenders to *see* one thing *as* another, or *act-as-if* one thing was another [3], [4], [8], [13], [17]. The paper will focus on one of the least intellectualist accounts of pretence proposed to date: Currie's [3] simulationist account of pretence.

Currie's ambition in advancing his account of pretence is clear: "The more we can account for what we do without supposing we need to think about doing it, the better" [3, 191]. Towards that end he regards pretence as involving imaginative transformations that, while not meta-representational, require *decentring*, or the "capacity to view the world from another perspective" [3, 211]. *Decentring* entails *off-line* simulation of possible scenarios, as will be discussed below.

However, even less intellectualist account of object-substitution pretend play is possible by applying O'Regan and Noë's [18] sensorimotor theory of perception (henceforth: SMTP), which will be argued for in the following way. Section 2 of the paper describes Currie's position on pretending in terms of *seeing-as*, followed by how the notion of seeing-as has been utilized by SMTP. How SMTP might contribute to pretending is clarified in Section 3. The suggestion here is that SMTP aids in understanding basic pretence by applying to one crucial aspect of Currie's notion of pretence – the perceptual *seeing-in* (or *experiencing-in*) – but reinterpreting it along more enactive lines (in terms of *seeing-affordances-in*). It will be shown that this is an important corrective to Currie, but it is not enough. For Currie holds that seeing-in is at most necessary but not sufficient for pretence, insisting that representational *seeing-as* understood as the capacity to make imaginative transformations is required even for the simplest acts of pretence. Section 4 discusses why *decentring* is deemed by Currie to be necessary to play the role of such transformations. Section 5 suggests possible ways to rebut against such arguments, indicating how it might be possible to account for the sorts of imaginative transformations that Currie thinks *seeing-as* requires, while bypassing the need for *off-line decentring* at its core. Section 6 proposes a positive account, where, in the place of *decentring*, it is argued that the relevant work might be done by *sensory imaginings* (understood as *on-line* perceptual activities), which are augmented by sensorimotor skills, certain understanding of Gibson's theory of affordances [7] and Merleau-Ponty's [14] account of directly perceived meaning, and further developed through narrative practices. This novel position aims to provide a first step toward a theory of basic pretence based on action rather than representation.

2 *Seeing-as* in Pretence and SMTP

What is required for *seeing* the banana *as* a phone? Currie explains that "in pretense, a creature may respond to the environment, but as it is transformed by imagination" [4, 275]. To *see-as*, "the pretending creature represents the world, not as it is, but as it might be" [4, 276), or *decentres*. In earlier work, Currie [3, 211] re-describes *decentring* in terms of a representational shift of perspectives:

> Decentring is the capacity to view the world from another perspective: to view the world as it was for me yesterday, as it is for you now, as it might be for me tomorrow, as it is according to some story. Decentring indicates the (relative) freedom from environmental constraint and sensitivity to representational content we think of as part of rationality.

Minimally, such imaginative transformations are needed to account for having controlled experiences in the absence of appropriate stimuli. As Harris and Kavanaugh claim, "pretence is similar to false belief in that actions stemming from both mental states are directed at situations that do not actually obtain" [8]. Thus, decentring is said to be necessary to allow the basic 'as-if' response to the environment, which does not produce direct stimuli.

One way of developing the notion of decentring might be to appeal to Recerative Imagination [5].[2] According to Currie and Ravnscroft, Recreative Imagination is simulation of perception, which, in turn, involves representing and manipulating perceptual contents *off-line*, or to "substitute one thought *content* for another" [5, 140]. It grants us the "ability to experience or think about the world from a perspective different than the one that experience presents" [5, 9]. If we understand Recreative Imagination as an empirical hypothesis about a mechanism involving simulation of perception, it might explain how decentring is done.

Pretence, to Currie, seems to be a higher cognitive activity after all. For example, Currie says that "there is a sense in which pretence *is* a 'higher' mental process … The child who pretends that Pig is dirty needs to have the first-order thought 'the Pig is dirty'; … this thought is tokened as part of an act of decentring" [3, 219]. Currie also claims that *seeing-as* involves "acting under a suppositional mode", where one can "consider an idea draw consequences from it, consider the evidence for it, and compare it with other ideas" [3, 233], or that pretence stands in a relationship to the act of pretending (understood as a mental state of imagining) like truth stands to believing [3, 205]. Thus, while Currie's account of pretence is to date the least intellectualist one, it is not clear that it goes as far as it ought to.[3]

In SMTP, in turn, we find an enactive understanding of *seeing-as*. The big idea behind SMTP is the stress it lays on the role of embodied activity over thought. O'Regan and Noë are the promoters of a view of perception that is intimately linked with action; they follow the motto that *perceiving is something we do* [18]. For example, Noë claims that seeing a cube as a cube is a form of embodied activity:

---

[2] To apply it to the banana-phone case, Currie and Ravenscroft [5, 33] claim: "(A) child holding a banana to her ear and speaking into it is pretending to make a telephone call when the behavior is accompanied, or perhaps driven by, the imagining that this thing, actually a banana, is a telephone."

[3] Decentring is presented as necessary to be *thinking* about the world [4, 277]. In the end, it does not seem for Leslie and Currie to be far off from each other, as the research question of Leslie's "What allows children to think of a banana as a phone?' [11] and Currie's 'What allows children to act as if the banana was a phone?" [3] could amount to the same thing.

> When you experience something as cubical, you experience it as presenting a definite sensorimotor profile. That is, you experience it as something whose appearance would vary in precise ways as you move in relation to it, or as it moves in relation to you [16, 117].

Although Noë thinks you cannot perceive without content (according to him, perception is both content and concept involving), he has an unusual story about perception that makes it an activity [16]. Thus, seeing the cube as a cube would, arguably, invoke an exercise of a basic sensorimotor skill that treats mere *seeing* as kind of *doing*.

Still, we do not have an account of how to apply SMTP to pretence. Can it be extended to illuminate our understanding of pretence? My suggestion is that if there is a natural connection between perception and imagination [5], then there is a potential contribution SMTP could make to the topic of pretence. Yet, even if we grant the validity of SMTP as applied to perception, it seems to, at best, only target the capacity to see X as X (e.g., a cube as a cube), but not X as Y (e.g., a banana as a phone). If pretence is supposed to be directed at that which is not perceptually present, how is SMTP relevant for pretence? I suggest that Currie's notion of perceptual capacity to *see-in* may fulfill this demand. The next section will first elaborate on Currie's notion of *seeing-in* that underlies the capacity to *see-as*. In agreement with Currie, I will claim that *seeing-in* may be a crucial capacity for engaging in basic pretence. But while Currie's notion is passive, this section will suggest that SMTP allows for a new understanding of *seeing-in* as an activity. Understanding *seeing-in* as playing an active role in *seeing-as* is crucial for further analysis of seeing-as.

3 *Seeing-in* and *Seeing-affordances-in*

To Currie [3], *seeing-in* (or *experiencing-in*) is a perceptual basis for *seeing-as*. It is important to pretence as it plays various enabling roles in pretence.[4] *Seeing-in* may be a precursor to *seeing-as*.[5] It is a phenomenon that occurs when one, for example, sees a woman in a picture or a face in the clouds. Currie claims,

> Such seeing-in does not involve seeing a woman, nor does it involve the perceptual illusion of seeing one; neither is it a case merely of judging that the picture represents a woman: it is genuinely perceptual phenomenon [3, 220].

Currie contends that this ability extends from seeing things in static objects to seeing things in human behaviours:

---

[4] "Seeing-in may constitute part of primitive basis of pretence, it enables pretence to be enacted and communicated without the necessity for full-blown conceptual capacities" [3, 222].

[5] Otherwise, presumably, one could transform anything into anything else. Seeing-in may, thus, be a kind of weak constraint that structures what is being imaginatively transformed.

> The next step (…) is to suggest that, just as we can see things in pictures, we can see things in simple mimetic acts. When someone moves in a certain way we can see in their movements such acts as driving a car, hitting a cricket ball, or nursing a baby (…). The movements might be exaggerated or stylized, but we can still see the action in the performance, just as we see a well-known face in its caricature [3, 221-222].

Given there is no woman or actual cricket game present to see, and there is no confusion (or illusion) as to what is occurring, it is a fair question to ask whether seeing-in is a genuine perceptual phenomenon after all. Bracketing that concern, and allowing for the sake of argument that it could be a perceptual phenomenon, the question I want to focus on is: could seeing-in be applied to pretend play cases without involving making judgements or inferencing? Invoking affordances could provide the basis for a more de-intellectualized account.

Crucial to the account proposed in this paper is the idea that seeing-in could be understood as *seeing-affordances-in*. [6] Drawing on Noë's account of Gibsonian affordances [16], I will propose that *seeing-affordances-in* is an activity (as opposed to passive thought process), in line with SMTP, that allows bypassing the need for off-line *decentring*.

Noë describes affordances in the following manner:

> Things in the environment, and properties of the environment, offer or afford the animal opportunities to do things (find shelter, climb up, hide under, etc.). (…) When you see a tree, you not only directly perceive a tree, but you directly perceive something up which you can climb. Gibson took this feature of his theory to be quite radical, for it suggested that we directly perceive meaning and value in the world; we do not impose meaning and value on the world [16, 105].

Noë's environmental affordances are best understood as possibilities for actions. The role of affordances is especially promising because it allows the possibility to see in acitivities, not just entities: we could be seeing in the 'banana-at-the-ear play' an affordance to play 'calling'. But for SMTP to secure the claim that banana affords 'telephone' play, what may be additionally needed is an account of social affordances in play, as will be proposed in further sections. The claim will be that what the banana affords will become fully meaningful when children get immersed in the intersubjective environment and interactions.

[6] Currie also speaks of seeing-affordances-in when referring to Millikan: "Millikan emphasizes the role of looking for and seeing affordances in the environment. (…) it seems to me that this may be a kind of seeing-in" [3, 220]. However, for Currie, the capacity to *see-affordances-in* plays a different role (the role of recognizing the pretence in others) than the one required for the sensorimotor account of pretence (enabling treating one object as another, or *seeing-as*). Thereby, it will be treated as a different concept.

4 Potential Worries

There are valid worries that may be raised by the standard representationalist approaches to pretence, which I will introduce in the present section. An account of seeing-in may not seem sufficient to explain pretence; after all, Currie himself insists that *seeing-as* understood as the capacity to make imaginative transformations is required even for the simplest acts of pretence. What follows are some reasons for thinking why we might need decentring.

Firstly, as mentioned earlier, such representational faculties are allegedly needed for one to "be directed at situations that do not actually obtain" [8] or to "stand back, cognitively speaking, from the immediate environment" [4, 275], for fulfillment of which Currie endorses making off-line simulations. With *seeing-affordances-in*, we also need to explain the possibility of seeing an action (calling) in the object (the banana). The problem is that there is no affordance to ring someone on the banana or dial a number, so how can objects afford special 'phone' actions, such as, e.g., calling or dialing? The worry here is that the question may have shifted from 'how is it possible to see objects (that are not there)' to 'how it is possible to see affordances (possibilities for actions), which are not there'.

Secondly, Currie claims that *seeing-in* constrains possibilities of actions, while in pretence, we have many possibilities of play. Seeing-in may not be enough because it has further constraints, such as being "fast, mandatory, encapsulated, very little dependent on learning…Try *not* seeing a person in the picture next time you look at a painted portrait," says Currie [3, 220-221].[7] The suggestion is that *seeing-in* perhaps would leave us with a very little room for voluntariness and creative choice in how or what we pretend.

Finally, an objection to an account of pretence involving *direct* perception is that it is supposed to be the *meaning* (not the environment) that guides pretence. According to Currie, "Vygotsky recognized (that) pretence is a form of decentring: the pretending creature is guided 'not only by immediate perception … but by … meaning'" [3, 211]. Meanings, traditionally, are understood as ideas or thoughts imposed on reality. The intellectualist assumption is that without representing the meaning of what is to be acted out, one could not get engaged in pretence in the first place. Thus, the direction of fit is supposed to be meaning to environment (adding new meaning 'phone' to the banana to pretend play 'banana-phone'), and not environment to meaning, which is what the seeing-affordances-in would propose.

[7] Similarly, it may be claimed that affordances are limiting. As Vygotsky claims, "(Things) dictate to the child what he must do: a door demands to be opened and closed, a staircase to be run up, a bell to be rung. In short, things have an inherent motivating force. (…) In play, things lose their motivating force. The child sees one thing but acts differently in relation to what he sees. Thus, a situation is reached in which the child begins to act independently of what he sees" [22, 11].

5 Rebuttal

This section will start with answering the three objections and will be followed by a suggestion of an account of pretence as inspired by SMTPs.

5.1 Affordances Are Present, Not Absent

The first objection is that one requires representations to refer to something that is absent. My rebuttal does not deal with the assumption that when objects are absent, representations are needed.[8] Rather, it deals with the assumption that in pretend play, we encounter absence, and we *must* stand back from the immediate environment. It is not clear whether all pretend play deals with such absence from the immediate environment. It is questionable whether in the situation when the child acts upon a prop (like in the banana-phone game), he or she ever acts independently of what is seen. That is because the banana, as an object, is part of the immediately present world, which affords acting upon. We may also think of affordances not only as properties of objects [16] but also a relational quality [1]. So while 'phoneness' property or 'buttons' property is absent in a banana, the shape of the banana is present for 'calling' to a human child when the banana is placed to an ear. Properties of objects as endorsed y Noë [16] as well as the history of past interactions as endorsed by Chemero [1] of the child shape how the banana is interacted with, which make 'calling' (or a way in which the object in question can be held, placed or turned around), in some sense, present. Thus, it is likely that in acting upon a prop (like in the banana-phone game), the player does not act independently of what is seen, but is guided by the prop and perceives in action what the prop affords. Acting upon affordances, which allows manipulating the possibilities of what objects or situations afford, answers the question of how it is possible to see something other than what is perceptually present.

5.2 Affordances Structure, Not Constrain

With seeing-affordances-in, we are not limited to see one way of interacting with an object. A banana may afford various actions in the context of play, such as playing 'phone' when held to an ear, playing 'hat' when held on the head or playing 'gun' when pointing it at someone. Yet, importantly, it is not the case that 'everything goes'; objects can also resist other kinds of play (e.g., playing 'human shoe' with a banana would be tough as the banana would get squashed). Thus, object affordances give novel possibilities of play, structuring play but not limiting possibilities. That objects can be played in more than one way is a view supported in recent

[8] Noë himself suggests that we need representations when the world is not immediately present. "Surely we sometimes need to think about the world in the world's absence (when it's dark, say, or when we're blind, or not at the location we're interested in), and for such purposes we must (in some sense) represent the world in thought" [16, 22].

psychological findings [15]. Objects are defined by their social and communicative uses, not their 'inherent meanings'.

Importantly, it is not the case that anything affords anything; that would not be very telling. For example, the banana does not afford as many things to a pro improviser, as opposed to a novice, as opposed to non-human animals without thumbs that could not lift it. An actor engaging in improvisations with objects has training, and such history of past interactions also shape the number of possibilities the object affords to one. Thus, looking at the individual capacities and history of engagements, not only perceiving entities, is important for this account [19]. For a person involved in social practices, something more can be afforded. This idea will be elaborated on in the next section.

5.3 Emergent Meaning and the Direction of Fit

With regard to the question of meaning, Vygotsky may have been right in noticing that some forms of pretence are framed by meaning, but his notion of meaning is not to be equated with decentring of the sort that Currie endorses. 'Meaning' in this context might include a wider grasp of active possibilities; it may be a different sense of seeing connections and possibilities for action. Vygotsky did not specify that meaning had to be representational or contentful, but notes that in pretend games young children are reliant on perceptually available information:

> Experiments and day-to-day observation clearly show that *it is impossible for very young children to separate the field of meaning from the visual field* because there is such intimate fusion between meaning and what is seen [22, 97].

Thus, there is a way to accommodate Vygotsky (but not Currie) in the claim that imaginative transformations are important for pretence, when Vygotsky's notion of meaning is not understood in terms of decentring but by applying an alternative conception of 'meaning'. One such alternative conception can be found in phenomenology, where the notion of meaning does not refer to mental contents, but to directly perceivable possibilities. According to Merleau-Ponty, perception is already meaningful; it "arouses the expectation of more than it contains, and … is therefore already charged with meaning" [14, 4]. The claim is that perception should not be opposed to imagination, even if they are in some way different, and that perception is already meaningful as it allows for novel possibilities to be perceived in the present object. Seeing a possibility is then anticipating something as happening. This is applicable to the pretend play in question, which can be explained not in terms of imposing new meanings in the form of rules for using the items on objects (such as imposing 'telephone' on a banana).[9] Instead, it can be accounted for with directly

[9] Even Vygotsky spoke of the lack of necessity of rules being established in advance of play: "The imaginary situation of any form of play already contains rules of behaviour, although it may not be a game with formulated rules laid down in advance" [22, 94].

seeing the meanings in terms of possibilities for action, or affordances (possibility of holding and using the banana as a phone).

There is a question to be asked about the direction of fit. With the notion of perceived affordances, the direction of fit seems to be from the environment (the object and what it affords) to meaning. This may be the problem with the notion of *seeing* affordances in objects, as the assumption would be that only those objects guide us in play. This can't be the full story. Some argue that object's uses are also not directly visible; their meaning lies in the public use [2]. The key move may be to stop thinking that objects have perceivable properties, and that what we may need are further capacities in order to pretend. The direction of fit from meaning to object can be preserved, with an adjustment. It is not the individual, representational capacity of decentring that affects the object, but in large part acting within intersubjective engagements that shape pretend play. That is, we needn't assume that in retreating to the meaning-environment direction of fit, we need to be decentring, but that something else (as meaning from intersubjective engagements) can play that role. The final section will gesture at how the intersubjective engagements shape pretend play.

6 The Positive Account and Its Benefits: Special Affordances Revisited

The positive account on offer here is that imaginative play is based on a strong link between perception and action, as proposed by O'Regan and Noë's SMTP [18]. SMTP highlights the role of sensorimotor engagements in creating the various possibilities for action, and constructing new meanings in intersubjective space. The seeing of possibilities for action may be extended by the application of know-how (or sensorimotor contingencies), as in the example of the 'pretend drying' of a toy elephant that got 'pretend wet' [12]. The 'drying' is a classic case of an outward behaviour that seemingly can only be explained by the theory that one carries out actual inferences 'behind the scenes' (such as: "if the pretend-water is being poured on the elephant then the elephant will be pretend-wet"). An alternative view of this situation is that one sees connections between the pouring movement in the case of the toy-elephant and, e.g., water being poured from the teapot, water poured on the child during a bath, or what showers are from stories and cartoons. The 'drying of wet things' behaviour may stem from applying what has been experienced in everyday 'taking-a-bath' contexts or narratives, to the 'wet-elephant' immediate play context. That is an example of applying know-how from perceptual activities stemming from the actor's history of past engagements to pretence.

That meanings can be constructed in intersubjective space is also suggested in De Jaegher and Di Paolo's [6] description of pretence in charades, where, even if the pretender started with an initial premise and understood the meaning of what's to be pretended, the way of acting out that the word on the card changes due to the breakdowns of social communication. In such cases, a 'shared meaning' of what is played emerges when the pretenders adjust their performance to the audience's needs, and audience's responses guide further changes in 'depicting' the meaning of the word on the card. De Jaegher and Di Paolo call this 'participatory sense-making' [6]. They suggest that novel meanings are established from mutual understanding, not from manifesting the initial premise.[10] The players' gestures change and evolve through patterns of coordination and breakdowns in the social setting.[11]

Social affordances might also take up some of the work in pretence that individual representations were thought to do [2]. The pretence may get started by imitating another person without knowing what the goal of the game is or following a clear script, which shows that an 'initial premise' with specific contents and rules of the game represented is not necessary. Imitating may be one way that children are brought into games – socially – without having to have all of the individual resources within themselves to frame the activity [19].

The connection between intersubjectivity and affordances is tight, which is apparent in Gibson's claim that affordances are neither in the environment, nor in the agent, but in the interactions, set up through a history of interactions (which leaves a possibility of their shifting) [7]. As Chemero in [1, 145] says,

> A better way to understand abilities [than to understand them as dispositions – *added comment*] is as functions. Functions depend on an individual animal's developmental history or the evolutionary history of the species, both of which occur in the context of the environment. Given this, it is actually more appropriate to understand abilities, like affordances, as being inherent not in animals, but in animal-environment systems. That is, like affordances, abilities are relations.

Thus, it is not the case that in the *object* one finds 'stable' affordances (e.g., "phoneness" of the banana shape); nor is it the case that it is something in the *individual* that allows these pretend actions (e.g., the representation of "phone" with a set meaning stored in the head and imposed on the banana), but a dynamic relationship between the agent and the object, which changes with movement, and novel interactions with other people and the objects, affording thereby novel ways of acting. How we develop the interactions with relevant affordances is to be discussed in future research.

---

[10] De Jaegher and Di Paolo [6] claim that the content of the intention determines the game, but that due to participatory sense making, shared meanings are developed so it is the sensitivity to others' understanding of what the pretender is projecting that is guiding his/her behavior. If that is right, then my prediction is that if hopping around does not do the job of conveying 'rabbit', the child will try other activities, stemming from learned stereotypical ways of playing.

[11] Hutto and Myin [10, 173] also claim that meanings can emerge from social interactions and are created in shared practices: "(The) very possibility of conceptual meaning, error and assessment requires an inter-subjective space. (…) Acquisition of such conceptual abilities depends on being able to have and share basic experiences with others." While they are talking about concepts in particular, it is plausible that their views can extend to pretend play affordances.

7 Conclusion

In this paper I have proposed an account of pretence as involving active seeing-in. I have argued that sensory imaginings augmented by sensorimotor skills suffice for playing the role of imaginative transformations required for *seeing-as* to explain the basic type of pretend play, like the banana-phone play. Without the need to represent what they're doing, children engage in pretend play by enacting typical routines stemming from perception. An analysis of what is required for imaginative transformations in *seeing-as* ceases to be an *off-line* mental activity and is instead understood as an *on-line* perceptual activity. The re-description of the notion of meaning from entertaining a set of rules to seeing possibilities of action allows to treat meanings as directly perceptible, and applicable at least in pretend games with props such as object-substitution play. S*eeing-affordances-in* plays a central role for pretence as it is an account of how objects and situations may be engaged with, using a basic capacity such as a sensorimotor mechanism. This novel approach to pretence suggests that pretence can be conceived as a way of acting that relies solely on embodied, perceptual and intersubjective skills, which is in line with O'Regan and Noë's view of perception: just as they claim that perceiving is a way of acting, so I claim that basic pretending is.

There are several advantages of this novel account of pretence as inspired by SMTP, particularly by its focus on activity and the role of affordances in perception. First, it is a candidate for the least demanding (and anti-intellectualist) view of pretence as per arguments from ecological validity, Ockham's razor or Morgan's canon. Yet, this attempt at a sensorimotor account of pretence may not extend to all forms of pretence for fictional activity. Undeniably there are more complex ways of pretending possible and they may invoke representing or making inferences, especially once full-blown language capacities are at play.[12] However, to explain pretence of 18-month-old children, an account of perception and action suffices. In line with known empirical studies, children at the age of 18 months engage in pretend games are reliant on perceptually available information [8], [13]. Moreover, mere seeing-in is not enough; action is necessary. The acts of pretend play are the parade cases of pretence. As Vygotsky observes, "child's play is imagination in action … we can say that imagination in adolescents and school children is play without action" [22, 93]. Presumably, just seeing one thing as another without acting comes later [20]. Thus, when dealing with the banana, the features of the object in line with what children see done with phones (their history of engagements) and their enacting of observed affordances best explains why playing with it as a phone is a natural thing to do. Hence, it is likely that the banana-phone object substitution play can reduce to utilization of solely such sensorimotor abilities, at the same time being a paradigm

[12] Just perceiving and acting may not be enough for a full-blown account of pretence; there are other aspects of pretence crucial for its success, such as imitating, responding to emotions, gesturing, smiling, giggling, context-sensitivity, and linguistic skills, to name but a few. Perceiving and acting are mentioned solely to suggest what minimally plays the role of imaginative transformations that underlie the banana-phone pretence.

case of basic pretence that more complex accounts of pretence can build on. As the burden was on the embodied and intersubjective theories to show whether any form of genuine pretence is possible without representing [21, 130], this account has shown that there is a possible space for non-representational pretence that applies to basic pretence, laying groundwork for its application to further types of pretence for future research.

Acknowledgements. This work was funded by the Marie-Curie Initial Training Network TESIS: "Towards an Embodied Science of InterSubjectivity" (FP7-PEOPLE-2010-ITN, 264828). I would like to thank Dan Hutto, Sam Coleman, Shaun Gallagher, Erik Myin, Dan Zahavi and the colleagues from the Centre for Subjectivity Research in Copenhagen for support received in relation to the work presented in the paper.

References

1. Chemero, A.: Radical Embodied Cognitive Science. MIT Press, Cambridge (2009)
2. Costall, A., Dreier, O.: Doing Things with Things: The Design and Use of Everyday Objects. Irish Journal of Sociology 17(2), 791–6035 (2006) ISSN 0791-6035
3. Currie, G.: Arts and Minds. Oxford University Press, Oxford (2004)
4. Currie, G.: Rationality, decentring, and the evidence for pretence in non-human animals. In: Hurley, S., Nudds, M. (eds.) Rational Animals? Oxford University Press, Oxford (2006)
5. Currie, G., Ravenscroft, I.: Recreative Minds: Imagination in Philosophy and Psychology. Oxford University Press, New York (2002)
6. De Jaegher, H., Di Paolo, E.: Participatory Sense-Making: An enactive approach to social cognition. Phenomenology and the Cognitive Sciences 6(4), 485–507 (2007)
7. Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston (1979)
8. Harris, P., Kavanaugh, D.: Young Children's Understanding of Pretense. Society for Research in Child Development, Monograph, vol. 58. University of Chicago Press, Chicago (1993)
9. Huttenlocher, J., Higgins, E.T.: Issues in the study of symbolic development. In: Collins, W. (ed.) Minnesota Symposia on Child Psychology, vol. 11, pp. 98–140. Erlbaum, Hillsdale (1978)
10. Hutto, D.D., Myin, E.: Radicalizing Enactivism: Basic Minds without Content. MIT Press, Cambridge (2013)
11. Leslie, A.: Pretense and representation: The origins of "theory of mind". Psychological Review 94, 412–426 (1987)
12. Liao, S., Gendler, T.: Pretence and Imagination. WIREs Cognitive Science (online publication) (2010)
13. Lillard, A.: Making sense of pretense. In: Lewis, C., Mitchell, P. (eds.) Children's Early Understanding of Mind: Origins and Development. Lawrence Erlbaum Associates, Hillsdale (1994)
14. Merleau-Ponty, M.: The Phenomenology of Perception. Routledge and Kegan Paul, London (1962)

15. Moro, C., Rodrıguez, C.: Production of signs and meaning-making process in triadic interaction at the prelinguistic level. A task for sociocultural analysis. The case of ostension. In: Abbey, E., Diriwachter, R. (eds.) Innovative Genesis: Microgenesis and the Constructive Mind, pp. 205–225. Information Age Publishing (2008)
16. Noë, A.: Action in Perception. MIT Press, Cambridge (2004)
17. Nichols, S., Stich, S.: Mindreading: An Integrated Account of Pretence, Self-Awarenessand Understanding of Other Minds. Oxford University Press, Oxford (2003)
18. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24(5), 939–1031 (2001)
19. Rucinska, Z.: Pretence as Engagement in Unscripted Routines. Pending in Phenomenology and Cognitive Sciences, special issue on Narrativity
20. Sainsbury, R.M.: Fiction and Fictionalism. Routledge, London (2009)
21. Spaulding, S.: Embodied Cognition and Mindreading. Mind and Language 25, 119–140 (2010)
22. Vygotsky, L.S.: Mind in Society: The Development of Higher Psychological Processes. Harvard University Press, Mass (1978)

Investigating Sensorimotor Contingencies in the Enactive Interface

Janet K. Gibbs and Kate Devlin

Goldsmiths, University of London. London, UK
k.devlin@gold.ac.uk

Abstract. Sensorimotor theories of perception have been widely investigated in the context of the perceiver's normal environment, but not in the context of virtual environments. There are clearly identified differences between perception of pictures and that of a real-world environment, but these differences have not been studied in the light of sensorimotor theory. Nagel et al.'s studies of sensory augmentation included a trial of their feelSpace belt in a computer-game environment, but with inconclusive results. We propose that the sensorimotor contingencies that apply in the context of a virtual environment are significantly different from those in the 'real world', and might account for the differences found between 'normal' and picture perception. Building on Froese et al.'s work on Enactive Interfaces, and on Visell's structure for sensory substitution, we consider how interfacing a sensory augmentation device with a computer game environment might provide the basis for fruitful research in this area.

1 Sensory Substitution, Augmentation and the Enactive Interface

In a paper of 2011, Froese et al. [1] propose a definition of an Enactive Interface (EI) as "a technological interface that is defined for the purpose of augmented sense-making". The authors see the interaction of sensory 'input' and active 'output' as two facets of the process of sense-making: "the activity by which an autonomous and adaptive agent maintains a meaningful relationship with its environment". Such sense-making emerges from goal-directed activity, not just as responses to stimuli. Thus an EI can be seen as "any piece of technology that is designed to permit its user to engage in additional modes of sense-making by enabling the goal-directed regulation of previously unavailable sensorimotor contingencies." Such devices are 'experientially transparent' ('ready-to-hand' in Heidegger's terminology), as compared with the 'cognitivist' approach to technological interface, in which "the user is forced to shift their attention to the abstract output of the device and must reason about what this output means for the course of action ... rather than being implicitly facilitated in perceiving what to do ...".

For Froese et al., the Enactive Interface encompasses what have been widely discussed as Sensory Substitution (SS) and Sensory Augmentation (SA), along

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_13,

with systems such as haptic interfaces, prosthetic and assistive devices, virtual reality (VR) systems, and many aspects of everyday human-computer interaction (HCI)[1].

Schmidmaier [2] gives an extensive overview of SS and SA systems, their technology, and the insights they provide into the workings of human perception albeit with a strong ‘passive perception’ flavour. He categorises historical and currently available systems according to the modality by which information is displayed to the user, and the modality of its source data. Thus, for example, Bach-y-Rita’s TVSS is categorised as Tactile/Visual, whilst Nagel’s feelSpace belt is Tactile/Spatial Awareness. EI devices need not be confined to ‘high-tech’ systems that most current research addresses. The blind person’s white stick is recognised as an example of Haptic+Audio/Spatial Awareness SS [3]. Some authors include the Braille system of writing, and Sign Language for those with auditory impairment, as further examples [2,3]; some would go so far as to include writing itself as an Auditory/Visual EI [4](cited [2]). By a similar argument, a drawing or photograph can be understood as a Visual/Visual EI; an audio player as Auditory/Auditory, and cinema or TV as a combination of both. For any such device to qualify as EI, we should require it to “implicitly facilitate perception”. For example, Nagel et al. [5] consider that ‘subcognitive processing’ occurs with their feelSpace device when the user is able to benefit from the signals provided without the expense of attentional resources. This is in line with Froese’s distinction between EI and the ‘cognitivist’ view of technological interface as cited above.

Visell [3] presents a somewhat different review of tactile SS, relating it both to our understanding of human perception and to the development of Human Machine Interaction (HMI). He presents the following useful structure for SS:

> “Information about the environment is typically acquired from sensors corresponding to modality A, and the information is transduced into a set of signals x(t) that are subsequently digitized. The sensors can be physical devices or they may correspond to measurements in a virtual environment. A coupling device maps the sensed data x onto a set of signals y(t) for driving the actuators of the display. The actuated display presents the information to a human sensory modality B, which is eventually transduced and processed by the intrinsic sensory system of the body.” (See Fig 1 below.)

Visell also notes: “One feature . . . that many have argued is crucial to the effectiveness of such systems is that the interaction loop is closed, through the affordance of user control over the position and orientation of the sensors, represented by the dashed lines in the figure.” It is important not to see Visell’s structure as relating only to the input side of an input/output (sensing/action)

[1] For the purposes of this paper, we shall use EI in this encompassing sense, but will continue to refer to SS, SA etc where the context requires the more specific reference–for example, when discussing a paper or proposition that relates only to the narrower concept(s).

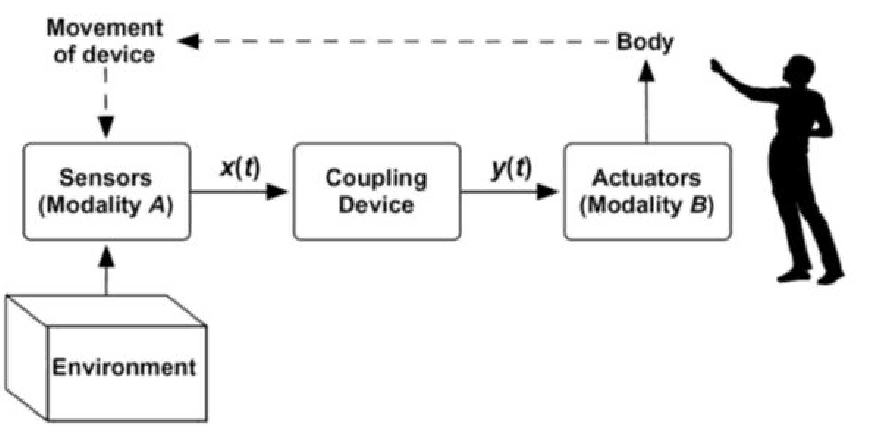

Fig. 1. Structure of a Sensory Substitution system (reproduced, with permission, from Yon Visell: Tactile Sensory Substitution: Models for Enaction in HCI [3])

system. 'User control' here should be understood to be an essential part of the user's active exploration of the environment – in EI terms, part of the process of active sense-making.

Visell's structure can readily be applied to the wider concept of EI, at least in the high tech examples. By specifying the mapping of digitized signals it would seem to exclude low tech systems; however, he also gives an alternative definition of SS as "the act of translating signals that are normally associated with sensory modality A to signals that can be detected via modality B". This definition applies equally well to the white stick as to its high-tech equivalent, the Enactive Torch (ET).

Table 1 shows how a range of EI systems would fit Visell's structure. As well as identifying the two modalities A and B, we have indicated where the sensors draw on a real or virtual environment, and distinguish between devices that aim to give experience of native and non-native modalities.[2]

2 Sensory Substitution and Augmentation in Perception Research

Sensory Substitution and Augmentation (SSA) is widely cited as offering support for Sensorimotor Theory of perception ([6,7,8,9,10]), as well as in more

[2] For this purpose we have categorised SS devices as delivering 'non-native' modalities: for example, the intended user of TVSS does not enjoy, as a native modality, the level of vision delivered by the device. However, we should note that there is no clear dividing line here, since blindfold sighted users, as well as users with a range of visual histories and of visual impairments, may all be included in trials of such devices.

Table 1. Examples of Enactive Interfaces, and how they fit into Visell's structure for SS devices

Coupling Device (Interface)	Nature of Environment (Real or Virtual)	Input to the EI (Modality A)	Actuators (Modality B)	Experienced Perceptual Modality	Control of Movement within Environment	Comments
Normal Perception	Real	All senses	All Senses	All native senses	Perceiver	Normal bodily movement
Still Picture	Virtual	Visual	Visual	Native Visual	None	
Moving Picture	Virtual	Visual	Visual	Native Visual	Device	Limited movement
Audio Soundtrack	Virtual	Auditory	Auditory	Native Auditory	Device	Limited movement
Interactive Game	Virtual	Visual, Auditory	Visual, Auditory	Native Visual, Auditory	Perceiver	Hand movements only, via mouse, keyboard, etc
Immersive Virtual Reality	Virtual	Multiple senses	Multiple senses	Multiple native senses	Perceiver	Full bodily movement
TVSS	Real	Visual	Tactile	Non-native visual	Perceiver	Limited movement (no eye movement)
Enactive Torch (also White Stick)	Real	Spatial	Tactile	Non-native spatial	Perceiver	Hand movements only
feelSpace	Real	Orientation	Tactile	Non-native orientation	Perceiver	Normal bodily movement
Computer assisted surgery	Real	Visual (Tactile?)	Visual (Tactile?)	Native Visual (Tactile?)	Perceiver	Normal bodily movement
Tactile Prosthesis	Real	Tactile	Tactile	Native Tactile	Perceiver	Normal bodily movement
feelSpace interfaced to Interactive Game	Virtual	Orientation	Tactile	Non-native orientation	Perceiver	Hand movements only, via mouse, keyboard, etc

conventional perception research (eg [2,3,10,11]). However, much experimental work in this area relies on uncontrolled small group or single case studies. Where controlled studies have been carried out, results have been more disappointing, and many speakers at a recent conference on SSA [10] concluded that evidence for phenomenal experience of the modality to be delivered by these devices just isn't there; some believing that such devices offer cognitive rather than perceptual experience[3]. On the other hand, it is notable that controlled studies (and, indeed, many small group studies) can measure their participants' use of the device under investigation in hours, usually over a study period of days or a few weeks ([1,5]); by contrast, what might be understood as genuine phenomenal experience of a delivered modality is repeatedly reported from the experience, often immersive, of single users over periods of months and years ([12,13]). It is also notable that, for visual SSA devices, there were considerable differences in the experiences obtained as between early-blind, late-blind and blindfold sighted users [10]. This raises the question as to how, and how much, each individual's previous perceptual experience influences the outcome of a trial. Again, there is significant difference between trials in the degree of directed training, as opposed to acclimatising experience, given to users ([1,5]). All of these factors make it difficult to draw any conclusions when comparing the outcomes of different studies.

This is clearly illustrated in Ward & Meijer's report [12] of the experience of a participant, PF, following immersive use of The vOICe auditory/visual SS device. PF, late blind at the age of 21, encountered The vOICe some 20 years later, and used it immersively from the age of 43. She reports that it took 3 months of immersive use "to learn enough so that I didn't have to consciously concentrate on it". Depth perception arose as "a kind of Eureka moment", after at least two months of "flat visual experiences of edges and shading". Five years later, she reported experiencing colour–which is not actually encoded by The vOICe: "Over time my brain seems to have developed, and pulled out everything it can from the soundscape and then used my memory to color everything". If a late blind user takes months to experience visual phenomenology, even with the benefit of remembered native experience, how much can we really learn from trials of only a few weeks?

3 Picture Perception – The Original Enactive Interface

We have argued above that pictures, in their various forms, satisfy the idea of an Enactive Interface. They also conform to Visell's structure for SS, which we have extended to apply to EI.

Just as Bach-y-Rita and Kercel saw writing as a kind of Visual/Auditory SS, so we might regard pictures as Visual/Visual EI. Froese et al. distinguish between devices that improve the function of an existing modality, such as spectacles for the short-sighted, and EI which should "enable the participant to generate and

[3] cf. Froese et al.'s distinction between 'cognitive' and 'enactive' technological interfaces.

make sense of qualitatively new forms of sensorimotor regularities". It might seem that this would not apply to a Visual/Visual device. However, Noë [6] has observed that:

> "Pictures construct partial environments. They actually contain perspectival properties such as apparent shapes and sizes, but they contain them not as projections from actual things, but as static elements. Pictures depict because they correspond to a reality of which, as perceivers, we have a sensorimotor grasp. Pictures are a very simple (in some senses of simple) kind of virtual space. What a picture and a depicted scene have in common is that they prompt us to draw on a common class of sensorimotor skills."

Note that we draw on a 'common' not an identical set of skills,. It is easy to see that some sensorimotor contingencies (SMCs) that apply in the real world, such as occlusion, parallax, looming etc, don't apply in a depicted image. In the more complex case of moving pictures (such as in film or in a first-person computer game), new contingencies come into play, as the images exhibit precisely these behaviours, but in relation to the camera or game character movement, rather than relative to any bodily movement of the observer. Furthermore, there is extensive research ([14]) into differences between 'normal' and picture perception–albeit from a conventional rather than a sensorimotor theoretical background.

Such differences include the duality of pictures, characterised variously as 'direct' and 'indirect' vision [15], 'twofoldness' [16] or 'conjoint representations' [17], whereby we see the picture both as an object in our natural environment, and its content as representing a different, virtual environment. These concepts of duality appear to have much in common with the distal and proximal perception of SSA; similarly, with Heidegger's three-fold concepts of 'ready-to-hand', 'present-at-hand' and 'unready-to-hand' in tool use; and with Froese et al.'s notion of transparency and opacity in the use of Enactive Interfaces. Cooper & Banks [18] and Sedgewick [19] discuss the distortions in depth perception that arise if we don't view a picture (in normal perspective) from its centre of projection. And Sedgewick [20] describes 'cross-talk' by which such distortions can be increased by emphasising features such as framing or surface texture that tell us we're looking at a picture.

And yet, there is a sense in which we do also draw on the contingencies of 'normal' vision, even when they don't apply in our picture. For example, in recognising a pictured object, we draw on the understanding of how its aspects (P-Properties, in Noë & O'Regan's terms) would change 'if I could move in relation to it within the depicted environment'. Thus the 'duality' of picture perception is, from a sensorimotor point of view, rather more complex. It seems that when we look at a picture, we may be exercising three different sets of sensorimotor skills–those applied to the picture as an object in our own environment; possibly some that apply only within the depicted environment; and those that would apply within the environment 'if I could move within it'.

It can be difficult to separate out these three sets. In the case of still pictures, it's not clear what contingencies could be specific to a depicted environment unless, for example, there are SMCs that account for the distortions of indirect vision. But in a moving (TV or cinema) image, we experience the contingencies arising from camera movement, even though we remain seated in one position. In the case of a first-person computer game, we have control of the character's movement, and therefore of the SMCs that arise–but not in the way we would in the real world, by walking around, turning our head etc; instead we control these contingencies by a very different set of hand movements to activate mouse and keyboard commands.

It is notable that we are very much unaware of all these differences. In discussing the importance of 'natural perspective' in our world, Hecht notes that much vision research is based on experiments using screen images, without any question as to their equivalence to 'everyday' visual perception.[4][14]. Pirenne [21] (cited [20]) suggests we may have an inbuilt mechanism to compensate for distortion due to looking at a picture from the wrong viewpoint–though Sedgewick queries whether such "a mechanism of formidable complexity" is likely to have evolved for the purpose of indirect perception alone. On the other hand, work with EI strongly suggests that whatever perceptual experience arises from these devices is a learned skill. In the developed world we have been looking at pictures (still or moving) almost as long as at the real world: perhaps it's just that these skills are equally transparent. With computer games, we may need to learn new manual skills, but these are applied to, and co-ordinated with, already transparent visual skills–though even here, there are some who experience 'visually induced motion sickness' (VIMS) ([22,23]). It might be instructive to look into responses to the new 3-D cinema technology, where there are clearly new contingencies to be learned, and where some viewers at least[5] find the experience distinctly uncomfortable.

4 Perception in Real and Virtual Space

We have seen that there are well-established differences between perception in the virtual space of still pictures and in our natural environment. It seems very likely that a similar range of differences apply to perception in more complex virtual environments, from moving pictures through computer games to immersive virtual reality. Research into picture perception has investigated these differences, primarily from the point of view of conventional perceptual theory, but no coherent picture has emerged as to why they arise. What new insight can sensorimotor theory of perception offer to this question?

We have suggested some ways in which visual SMCs do differ between particular virtual environments and the real world, and would predict that these should lead to phenomenal differences in perceptual experience. On the other hand,

[4] This practice may well arise from representational theories of perception: if visual perception is taken to be based on internal processing of a sequence of 'snapshot-like' retinal images–2-dimensional projections of 3-dimensional scenes–then it is easy to assume that seeing a snapshot is essentially the same as seeing its real world original.

[5] These authors included.

visual perception in both real and virtual environments is deeply ingrained in our experience, and the two are interrelated in such complex ways that it seems almost impossible to disentangle them. Looking at the broader field of Enactive Interfaces, it might be more fruitful to compare experiences that are both less well established, and in themselves less complex.

Table 2 offers a grouping of EI systems that suggests a way forward by making comparisons between categories on each dimension. Studies of picture perception have compared Native (visual) perception in Real and Virtual environments. SSA studies have focussed on Real environments, investigating the possibility of experiencing mainly visual and auditory perception as a Non-native, as compared with Native modality. However, two SSA devices in the Real/Non-native group relate to much less complex SMCs.

Table 2. Cross-tabulation of Enactive Interfaces on Real/Virtual vs Native/Non-Native dimensions

	Real (R)	Virtual (V)
Native (N)	Native perceptual experience of the real world–our normal mode of perception; also tactile prosthetics and systems such as computer-assisted surgery.	Native perceptual experience of a virtual environment, as exemplified by still and moving pictures, and by interactive computer games and immersive virtual reality.
Non-native (NN)	Non-native perceptual experience of the real world, as in SSA systems such as TVSS, feelSpace, ET and the white stick.	Non-native perceptual experience of virtual environment, as in Nagel et al.'s trial of their feelSpace device interfaced to a computer game.

Nagel et al., in their feelSpace trials, aim to deliver a non-native sense of spatial orientation[6]. Even though the team's later studies [25] suggest that users' phenomenal experience can be much more complex than simply knowing which way is north, the SMCs themselves are relatively straightforward: the tactile

[6] Awareness of directional orientation is regarded as a non-native perceptual modality in humans, but is native to some species of birds and animals who have a built-in ability to respond to the earth's magnetic field[5]. However Levinson [24] describes the aboriginal Guugu Yimithirr language group in Australia, who have no words in their language for egocentric concepts such as 'left', 'right', 'in front', 'behind', etc; instead, spatial references are made in terms of something very like our 'north', 'south', 'west' and 'east'. He notes that "GY speakers invariably seem to know, day and night, familiar or unfamiliar location, whether sitting still or traveling in a vehicle, where the cardinal directions lie." In the light of this extreme example, can we rule out the possibility that some individuals, whether consciously or not, can have a better 'sense of direction' than most, simply in response to everyday environmental cues?

sensation of the belt depends simply on the wearer's orientation about a vertical axis through the trunk. Froese et al.'s Enactive Torch (ET) also aims to deliver a form of spatial awareness, comparable to the blind person's white stick. SMCs here are less straightforward than feelSpace, but still much less rich than in visual perception, or even in Tactile/Visual or Auditory/Visual SS. Largely because of this relative simplicity, Froese et al. have proposed the ET as a "minimal EI" for the purpose of perceptual research [1].

Finally, in the Virtual/Non-native grouping, Nagel et al. have also reported on trials of using their feelSpace belt interfaced with a computer game environment [5]. Results were disappointing, but perhaps not surprising according to sensorimotor theory: differences in SMCs between the users' training in their Real environment and tests in the Virtual environment would readily account for this.

It would not be difficult to repeat Nagel et al.'s Virtual/Non-native trial using a suitably adapted feelSpace, ET or other comparable device; with training and tests all carried out in virtual environments delivering essentially the same SMCs. Assuming that some kind of subjective and/or measurable benefit is experienced in interacting with the virtual environment, the scene would then be set for further study in various directions. With suitable programming, variations in the environment, and in the way the device is used, can be investigated in a more controlled way than might be practicable in the Real world; users would have greater consistency in their non-experience of the proposed modality to be delivered; and the less complex nature of the modality might lead to shorter learning curves, which would make setting up controlled trials more practicable.

As well as studying a device and its usage within the Virtual/Non-native category, different EI systems may be compared both within and between categories. Such comparisons could shed light on the sort of problems we have discussed above. For example, differences and similarities between feelSpace in the real and virtual worlds (R/NN vs V/NN) may help us to understand the differences between direct and indirect vision (R/N vs V/N). Similarly, a comparison of feelSpace or ET experience against visual experience in the context of computer games (V/NN vs V/N) might suggest further lines of study as between SSA and normal perception (R/NN vs R/N). We offer some suggested lines of investigation using EI in a Virtual Environment, as follows:

Perceptual Experience

- Are there significant differences in experience as between using a system in the game environment and in the real world? (e.g. due to different SMCs arising from manipulating mouse/keyboard vs walking about)
- Is experience affected by delivery method: e.g. could feelSpace be worn (suitably scaled down) as a wrist strap, or as a Tongue Display Unit? Are skills transferable between these delivery methods?
- Are there differences in SMCs applicable to controlling first and third person characters? Would training in one context be immediately transferable to the

other? If not, what is the nature of the differences between them, and how would they effect subjective experience and/or measurable performance?
- Would we get the same sort of answers to the above questions for different types of device, such as feelSpace vs ET, in comparable game environments?

Training and Acclimatisation

- How much training and acclimatization are needed for benefits to be enjoyed? Is acclimatization sufficient without directed training?
- Are there factors in the virtual environment that might affect training requirements as compared with using the device in the real world? (e.g. is more or less training required? if so, why?)
- Are simpler modalities such as Spatial Orientation and Spatial Awareness easier to learn than, for example, visual/motor control of a character in the same environment?

5 Conclusion

The object of this paper was to explore the relationship between perception in the contexts of 'real' and 'virtual' environments, in the light of sensorimotor theories of perception, and particularly in the context of Enactive Interfaces; and to consider what sort of a device might be used to study this relationship empirically.

It is clear from the literature that research in areas such as picture perception has not yet taken much account of sensorimotor and related theories: the underlying assumptions of conventional 'snapshot' theories of perception are so well established that it is rarely found necessary even to mention them. Similarly, it appears that work done in picture perception has not always filtered through to research in EIs, although there is certainly some common ground. As a result, the unspoken assumption that images viewed on screen are processed in the same way as real world images is rarely challenged.

In attempting to pull together these different strands of research, we have proposed a classification of EI systems on two axes: whether the environment for perception mediated by an interface is 'real' or 'virtual'; and whether the perceptual modality delivered by the interface is native to the user or not. In the light of this classification, we have proposed that an SA device interfaced to a virtual computer game environment would offer scope for fruitful study. Such a device, classified as Virtual/Non-native, could be studied alongside comparable interfaces in other categories, such as the same SA device used in the real world (Real/Non-native), and the same computer game environment without the benefit of SA (Virtual/Native). It could be used both to pursue further work begun by others in the field, and to explore new avenues not yet studied. As a result, we anticipate that a valuable contribution can be made in drawing together research from a number of related areas whose work has so far tended to progress independently.

References

1. Froese, T., McGann, M., Bigge, W., Spiers, A., Seth, A.K.: The enactive torch: a new tool for the science of perception. IEEE Transactions on Haptics 5(4), 365–375 (2012)
2. Schmidmaier, M.: Sensory substitution systems. In: Media Informatics Advanced Seminar on Multimodal Human-Computer Interaction (2011)
3. Visell, Y.: Tactile sensory substitution: models for enaction in HCI. Interacting with Computers 21, 38–53 (2009)
4. Bach-y Rita, P., Kercel, S.J.: Sensory substitution and the human-machine interface. Trends in Cognitive Sciences 7(12) (December 2003)
5. Nagel, S., Carl, C., Kringe, T., Maertin, R., Koenig, P.: Beyond sensory substitution: flearning the sixth sense. Journal of Neural Engineering 2, R12–R26 (2005)
6. Noë, A.: Action in Perception. Representation and Mind. The MIT Press, Cambridge (2004)
7. O'Regan, J.K.: Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness. Oxford University Press (2011)
8. O'Regan, J., Noë, A.: What it is like to see: A sensorimotor theory of perceptual experience. Synthese 129(1), 79–103 (2001)
9. O'Regan, J., Noë, A., et al.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24(5), 939–972 (2001)
10. Sensory Substitution and Augmentation. Unpublished Conference organised by the Centre for the Study of Perceptual Experience and the Network for Sensory Research, London (March 2013)
11. Bach-y Rita, P.: Brain mechanisms in sensory substition. Academic Press, New York (1972)
12. Ward, J., Meijer, P.: Visual experiences in the blind induced by an auditory sensory substitution device. Consciousness & Cognition 19(1), 492–500 (2010)
13. Guarniero, G.: Tactile vision: a personal view. Journal of Visual Impairment & Blindness 71(3), 125–130 (1977)
14. Hecht, H., Schwartz, R., Atherton, M. (eds.): Looking into pictures: an interdisciplinary approach to pictorial space. The MIT Press, Cambridge (2003)
15. Gibson, J.J.: The ecological approach to visual perception. Lawrence Erlbaum Assoc., Hillsdale (1979)
16. Wollheim, R.: In defense of seeing-in. In: Hecht, H., Schwartz, R., Atherton, M. (eds.) Looking into Pictures: An Interdisciplinary Approach to Pictorial Space. The MIT Press, Cambridge (2003)
17. Mausfeld, R.: Conjoint representations and the mental capacity for multiple simultaneous perspectives. In: Hecht, H., Schwartz, R., Atherton, M. (eds.) Looking into Pictures: An Interdisciplinary Approach to Pictorial Space. The MIT Press, Cambridge (2003)
18. Cooper, E.A., Banks, M.: Perception of depth in pictures when viewed from the wrong distance. Journal of Vision 12(9) (August 2012)
19. Sedgewick, H.A.: The effects of viewpoint on the virtual space of pictures. In: Ellis, S.R. (ed.) Pictorial Communication in Virtual and Real Environments, pp. 460–479. Taylor & Francis, New York
20. Sedgewick, H.A.: Relating direct and indirect perception of spatial layout. In: Hecht, H., Schwartz, R., Atherton, M. (eds.) Looking into Pictures: An Interdisciplinary Approach to Pictorial Space. The MIT Press, Cambridge (2003)

21. Pirenne, M.: Optics, painting and photography. Cambridge University Press, Cambridge (1970)
22. Keshavarz, B., Hecht, H.: Visually induced motion sickness and presence in video games. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting, vol. 56, p. 1753 (2012)
23. Merhi, O.: Motion sickness, console video games and head-mounted displays. Human Factors: Journal of the Human Factors and Ergonomics Society 49(5), 920–935 (2007)
24. Levinson, S.C.: Language and cognition: the cognitive consequences of spatial description in guugu yimithirr. Journal of Linguistic Anthropology 7(1), 8–13 (1997)
25. Kaercher, S.M., Fenzlaff, S., Hartmann, D., Nagel, S.K., Koenig, P.: Sensory augmentation for the blind. In: Frontiers in Human Neuroscience (March 2012)

Non-representational Interaction Design

Marco Gillies[1] and Andrea Kleinsmith[2]

[1] Embodied Audio-Visual Interaction, Department of Computing
Goldsmiths, University of London, London, UK SE14 6NW
m.gillies@gold.ac.uk
http://www.doc.gold.ac.uk/~mas02mg/
[2] Virtual Experiences Research Group
Department of Computer and Information Science and Engineering
University of Florida

Abstract. This paper presents how non-representational views of cognition can inform interaction design as it moves from traditional graphical user interfaces to more bodily forms of interaction such as gesture or movement tracking. We argue that the true value of these "bodily" interfaces is that they can tap our prior skills for interacting in the world. However, these skills are highly non-representational and so traditional representational approaches to interaction design will fail to capture them effectively. We propose interactive machine learning as an alternative approach to interaction design that is able to capture non-representational sensori-motor couplings by allowing people to design by performing actions rather than by representing them. We present an example of this approach applied to designing interactions with video game characters.

Keywords: interaction design, bodily interaction, interactive machine learning.

Interaction Design is a discipline saturated with representations. The primary interaction mechanism with modern computers is a graphical interface which is composed largely of visual representations of the computer system. These can include a wide variety of displays representing anything from documents to scientific data sets to social interactions. The interface also includes many different representations of possible actions that people can take from simple buttons to directly manipulable visualisations. The implementations of the interfaces also require complex logical representations, as any feature of the world or computer system has to be represented in program code, normally with explicit representations of items in the world and explicit mapping of these representations to the visual interface.

These representations are in no way incompatible with a non-representational view of cognition. They are, after all, external representations, part of our distributed cognitive apparatus that we can manipulate to achieve our aims. They largely exist in domains where we have very well established forms of external representation that most educated people are well versed in manipulating as

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_14,

part of our working process: written language, diagrams and photographs. In fact, the graphical user interface is now so ubiquitous that it is a key part of the distributed cognitive apparatus of most adults in the developed world.

So we could be happy that a representational approach to interaction design and a non-representational approach to cognition are compatible if it were not for certain new developments in interaction design. New developments in human computer interaction are leaving the confines of the computer monitor and entering the 3 dimensional world of our full body movements. This paper will argue that this move will force interaction design to take account of non-representational cognition.

1 Bodily Interaction

The use of body movement to interact with computers has a long history, going back to pioneering work from the 1970s by researchers such as Myron Krueger[1]. However, only in recent years has it become feasible to create low cost, mass market bodily interfaces, due to the advent of commercial movement tracking devices such as the Microsoft Kinect™or the accelerometers built into modern smart phones. These new devices enable us to make use of large scale movements of many different parts of the body, in contrast to the small scale movements required by a keyboard and mouse. Much of the research in this areas has concerned gesture recognition and it's use in user interfaces, for example the work of Bevilacqua et al.[2] or Fails and Olsen[3]. However, there is also research that makes use of holistic movement in interaction ranging from dance controlled music (Antle et al. [4]) and expressive digital musical instruments (Fiebrink [5]) to interactive art (Snibbe [6]) and body monitoring for healthcare (Fergus [7]). With devices such as the Microsoft Kinect™body movement interfaces are now being used by ordinary consumers, with the first area of growth being in video games, given that most devices are marketed as video game controllers. However, their use is spreading, even if they remain attached to video game consoles, body movement interfaces are being used for other activities, for example the WiiFit software uses a games console as a platform for exercise and health.

2 What Is Natural about "Natural User Interfaces"?

What is the value of this for interaction? The type of interaction I have been describing has been marketed by Microsoft and others as "Natural User Interfaces": interfaces that are claimed to be so natural that they do not need to be learned. The logic behind this phrase is that, because body movements come naturally to us, a body movement interface will be natural. This idea has been criticised by many people, most notably by Norman in his article "Natural User Interfaces are not natural"[8] in which he argues that bodily interfaces can suffer from many problems associated with traditional interfaces (such as the difficulty of remembering gesture) as well as new problems (the ephemerality of gestures and lack of visual feedback). So is there value in the intuition that bodily interfaces are natural, and if so what is that value and why is it often not seen in existing interfaces?

I would argue that there is a fundamental difference in the nature of bodily interfaces and traditional interfaces. Jacob et al.[9] propose that a variety of new forms of interaction, including bodily interaction, are successful because they leverage a different set of our pre-existing skills from traditions GUIs. While a graphical user interface leverages our skills in manipulating external visual and symbolic representations, bodily interfaces leverage leverage skills related to body and environmental awareness. The skills that enable us to move and act in the world. Similarly, Dourish [10] proposes that we analyse interaction in terms of embodiment which he defines as: "the property of our engagement with the world that allows us to make it meaningful". This leads him to define Embodied Interaction as "the creation, manipulation, and sharing of meaning through engaged interaction with artefacts". While he applies this definition to both traditional and new forms of interaction, the nature of this engaged interaction is very different in bodily interfaces. Following Jacob we could say that, in a successful bodily interface, this engaged interaction can be the same form of engagement we have with our bodies and environment in our daily lives and we can therefore re-use our existing skills that enable us to engage with the world.

If we take a non-representational, sensorimotor view of perception and action these skills are very different from the skills of a traditional interface involving manipulation of representations. This view allows us to keep the intuition that bodily interfaces are different from graphical user interfaces and explain what is meant by natural in the phrase "natural user interface" (the so-called natural skills are non-representational sensorimotor skills), while also allowing us to be critical of the claims of bodily interfaces. Natural user interfaces, on this view, are only natural if they take account of the non-representational, sensorimotor nature of our body movement skills. Body movement interfaces which are just extensions of a symbolic, representational interface which are just a more physically tiring version of a GUI.

A good example of this is gestural interaction. A common implementation of this form of interface is to have a number of pre-defined gestures that can be mapped to actions in the interface. This is one of the types of interface that Norman[8] criticises. When done badly there is a fairly arbitrary mapping between a symbolic gesture and a symbolic action. Users' body movements are used as part of a representation manipulation task. There is nothing wrong with this *per se* but it does not live up to the hype of natural user interfaces and is not much different from a traditional GUI. In fact, as Norman notes, it can be worse, as users do not have a visual cue to remind them which gestures they should be performing. This makes it closer to a textual command line interface where users must remember obscure commands with no visual prompts. Gestural user interfaces do not have to be like this.

These problems can be avoided if we think of gestural interfaces as tapping sensorimotor skills, not representation manipulation skills. For example, the work of Bevilacqua et al.[2] uses gesture to control music. In this work, gestures are tracked continuously rather than being simply recognised at the end of the gesture. This allows users to continuously control the production of sound throughout the time they are performing the gesture, rather than triggering the gesture at the end. This seemingly simple difference transforms the task from representation manipulation (producing a symbolic gesture and expecting a discrete response) to a tight sensorimotor loop in which the auditory feedback can influence movement which in turn controls the audio. A more familiar example of this form of continuous feedback is the touch screen "pinch to zoom" gesture developed for the iPhone^{TM}. In this gesture an image resizes dynamically and continuously in response to the users' fingers moving together and apart. This continuous feedback and interaction enables a sensorimotor loop that can leverage our real world movement skills.

A second feature of Bevilacqua et al.'s[2] system is that is allows users to easily define their own gestures and the do so by acting out those gestures while listening to the music to be controlled. I will come back to this feature in more detail later, but for now we can note that it means that gestures are not limited to a set of pre-defined symbolic gestures. Users can define movements that feel natural to them for controlling a particular type of music. What does "natural" mean in this context? Again, it means that the user already has a learnt sensorimotor mapping between the music and a movement (for example a certain way of tapping the hands in response to a beat).

3 How Do We Design Non-representational Interactions?

This brings us to the question of how we design bodily interactions. Interaction design for traditional user interfaces is largely a task of designing representations, for example, the layout of widget or the display of information. If we take successful bodily interaction to require non-representational sensorimotor skills, then the task of designing it must be something very different. Rather than designing representations, it means designing sensorimotor couplings. In fact, designing sensori-motor couplings are not restricted to bodily interactions. The two are not as distinct as I have presented them here. Viewed from this perspective, direct manipulation GUIs involve elements of sensorimotor coupling and most bodily interfaces are likely to involve some external representations.

The key challenge is that, if we are designing sensorimotor couplings that we perform without mental or external representations, then we must design them without representations. If we do not use mental representations to perform these skills and we do not have a well established system of external representations to handle them, we are unlikely to be able to explicitly represent details of our movement skills. If we use traditional, representational methods of interaction design we will be asking designers to provide explicit representations of low level details of their movements, to which they have no conscious access. This is likely to be an impossible task. This is one explanation of the problems of existing bodily interaction techniques. Designers have to explicitly represents their movements in numerical terms programme code or other formal tools but they do not have this level of access to complex sensorimotor skills. So instead, they fall back on representing simple features of the movement (responding to position or velocity) or falling back on traditional user interface metaphors (buttons in 3D space). Neither of these approaches are designed to make effective use of sensorimotor skills, rather they are driven by what can easily be represented in code. The result is interactions that are not "natural" in the terms we have discussed above and which are hard to remember and perform.

We need a different approach. If we are to successfully create complex movement interfaces we must develop software tools that support creating based on our existing, non-representational sensorimotor skills. That means we cannot ask designers to form explicit representations of their movements. Instead we should allow designers to design interactions by directly applying their bodily skills. The should define movements by moving.

Designing by moving can be enabled by software tools that allow designers to specify interaction by giving examples of those interactions. This type of tool can be implemented using machine learning techniques to infer recognition models from these examples. However, traditional batch approaches to machine learning are not particularly well suited to a design process as designers must collect large amounts or data initially and then must rely on the algorithm to produce the desired results based on this data. This makes it hard to support the iterative processes and refinement that are key to successful interaction design. Interactive Machine Learning (Fails and Olsen [3]) is a approach to machine learning which attempts to overcome these issues by making user interaction central to the learning process. Users provide training data interactively and in doing so progressively refine the classifiers they are creating. This approach has been used in a variety of domains from image classification (Talbot et al. [11]) to end user training of spam filters (Kulesza et al. [12]). This method can allow designers to design interactions by interactively providing examples of movement which would be used to train and refine a machine learning algorithm that controls the interaction. This would make the process of designing one of performing examples of movement. Designers can design by performing movements that emerge out of their sensorimotor skills without ever having to form explicit, detailed representations of the movements they are performing. Fiebrink et al. [13] have applied interactive machine learning to gestural control of digital musical instruments. Their participants noted that interactive machine learning provided a fundamentally different way of designing which focused on direct embodied movement rather than analyzing gestures in terms of specific features.

Fig. 1. The actor being motion captured (left) and a participant interacting with the virtual character (right)

4 Non-representational Interaction Design for Video Game Characters

This section presents a prototype that attempts to use Interactive Machine Learning to enable the type of non-representational interaction design that we are proposing. It is a system for designing bodily interaction with video game characters. Two people can play the roles of the video game character and the player, showing how the character should respond by acting out the movements themselves. This allows them to design movements in a natural way, by moving, rather than having to think about explicit representations of their movement. The motion of both participants are recorded and synchronized. This data is then used as input to a machine learning algorithm which learns a model for automatically controlling a video game character so that it responds in the same way as the people designing it.

The capture set up enables two people to improvise interaction in order to design the virtual agent. One participant plays the part of the agent while the other plays the part of the player. The two participants are in different spaces and their movements are streamed live to each other. The movements of the performer playing the agent are motion captured using an Optitrack™optical motion capture system and mapped live onto the agent, which can be seen live by the other participant. This second participant's movements are recorded using a Microsoft Kinect™consumer motion tracking device. His or her movements are visible to the other participant as a live video stream. Both sets of data are synchronised and recorded as input to our machine learning system. For interaction with the AI agent, the participants' set up is identical; they interact with a virtual agent via the Kinect, but in this case the agent's actions are selected by the learnt model rather than being controlled live by a participant in motion capture.

Once data had been recorded the participants use a desktop computer to select clips from the recorded data that will be used as input to train the machine learning classifier. The resulting clips contain two types of data. Firstly they contain a range of motion capture data which can be used to animate the virtual agent in response to the player's action. Secondly they contain kinect data, which can be used to classify the player's actions. The kinect data at the start of the clip is taken to be the action that produced the response contained in the clip and this is used as input for training the machine learning engine, for which we used a decision tree algorithm.

The final result of our system is a virtual agent that responds live to a player's actions as sensed by the Microsoft Kinect^{TM}. The task of the decision tree is therefore to determine which action the participant had performed and then select a suitable response from the agent. The decision tree is used to classify the participant's posture into one of the possible actions. This classification is then used to select an animation clip to play back on the virtual agent.

We tested the system by working with a physical performer to design an interactive character. The performer was a professional who worked in the theatre and also taught performance, and whose practice centred on physical body movement. As such he had considerable expertise in movement and was well suited to designing bodily interaction. He worked with a number of members of the public to design the interaction and was encouraged to use his own working practices as far as possible. He used a long process of physical improvisation, which is a key part of his performance practice. This enabled him to design the detail of the movement. While he was readily able to talk at a high level about the movements and given them names, movement improvisation was key to working out the details of the movements and interactions. The prototype enabled him to design the interaction entirely through this process of physical improvisation without ever having to make explicit low level details of movement. Our non-representation approach therefore fitted easily with his practice as performer and movement expert.

5 Conclusion

Sensorimotor theory provides a powerful lens for understanding both what is compelling about bodily interaction and how it can fail. Seeing bodily interfaces as a way of leveraging our sensorimotor skills that are left out of traditional representational user interfaces helps explain the intuition that this type of interface is some how "natural" while also explaining how many actual example fail to be natural by falling back on representational gestures. More importantly, it shows us a way forward for how we can create interfaces that do make better use of our sensorimotor skills. Since our sensori-motor skills are highly non-representational, design approaches that rely on representing actions will never capture them apporpriately. Interactive machine learning on the other hand allows people to design interactions by performing actions, thus allowing them to directly make use of their non-representational skills. The examples presented in this paper provide an idea of what these new design method might be like, while the central argument is a call to make sure we do change the way we design interfaces to adapt to our non-representational skills.

References

1. Krueger, M.W.: Responsive environments. In: Proceedings of the National Computer Conference, AFIPS 1977, June 13-16, pp. 423–433. ACM, New York (1977)
2. Bevilacqua, F., Zamborlin, B., Sypniewski, A., Schnell, N., Guédy, F., Rasamimanana, N.: Continuous realtime gesture following and recognition. In: Kopp, S., Wachsmuth, I. (eds.) GW 2009. LNCS (LNAI), vol. 5934, pp. 73–84. Springer, Heidelberg (2010)
3. Fails, J.A., Olsen Jr., D.R.: Interactive machine learning. In: Proceedings of the 8th International Conference on Intelligent User Interfaces, IUI 2003, pp. 39–45. ACM, New York (2003)
4. Antle, A.N., Corness, G., Droumeva, M.: What the body knows: Exploring the benefits of embodied metaphors in hybrid physical digital environments. Interacting with Computers 21(1-2), 66–75 (2009)
5. Fiebrink, R., Cook, P.R., Trueman, D.: Human model evaluation in interactive supervised learning. In: Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems, CHI 2011, pp. 147–156. ACM, New York (2011)
6. Snibbe, S.S., Raffle, H.S.: Social immersive media: pursuing best practices for multi-user interactive camera/projector exhibits. In: Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, pp. 1447–1456. ACM, New York (2009)
7. Fergus, P., Haggerty, J., Taylor, M., Bracegirdle, L.: Towards a whole body sensing platform for healthcare applications. In: England, D. (ed.) Whole Body Interaction. Human-Computer Interaction Series, pp. 135–149. Springer, London (2011)
8. Norman, D.A.: Natural user interfaces are not natural. Interactions 17(3), 6–10 (2010)
9. Jacob, R.J., Girouard, A., Hirshfield, L.M., Horn, M.S., Shaer, O., Solovey, E.T., Zigelbaum, J.: Reality-based interaction: a framework for post-wimp interfaces. In: CHI 2008: Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems, pp. 201–210. ACM, New York (2008)
10. Dourish, P.: Where The Action Is: The Foundations of Embodied Interaction. MIT Press (2001)
11. Talbot, J., Lee, B., Kapoor, A., Tan, D.S.: Ensemblematrix: interactive visualization to support machine learning with multiple classifiers. In: CHI 2009: Proceedings of the 27th International Conference on Human Factors in Computing Systems, pp. 1283–1292. ACM, New York (2009)
12. Kulesza, T., Stumpf, S., Wong, W.K., Burnett, M.M., Perona, S., Ko, A., Oberst, I.: Why-oriented end-user debugging of naive bayes text classification. ACM Trans. Interact. 1(1), 2:1–2:31 (2011)
13. Fiebrink, R.: Real-time Human Interaction with Supervised Learning Algorithms for Music Composition and Performance. PhD thesis, Princeton University, Princeton, NJ, USA (January 2011)

Minimally Cognitive Robotics: Body Schema, Forward Models, and Sensorimotor Contingencies in a Quadruped Machine

Matej Hoffmann[1,2]

[1] Artificial Intelligence Laboratory, Department of Informatics, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland
[2] iCub Facility, Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genoa, Italy
matej.hoffmann@iit.it

Abstract. In response to the cognitivistic paradigm and its problems, the embodied cognition viewpoint was proposed. In robotics, this resulted in a radical move away from higher-level cognitive functions toward direct, almost "brain-less" interaction with the environment (e.g., behavior-based robotics). While some remarkable behaviors were demonstrated, the complexity of tasks the agents could master remained limited. A natural extension of this approach lies in letting the agents extract regularities in sensorimotor space and exploit them for more effective action guidance. We will use a collection of case studies featuring a quadruped robot to concretely explore this space of minimally cognitive phenomena and contrast the concepts of body schema, forward internal models, and sensorimotor contingencies. The results will be interpreted from a "grounded cognition" and a non-representationalist or enactive perspective. Finally, the utility of robots as cognitive science tools and their compatibility with different cognitive science paradigms will be discussed.

Keywords: developmental robotics, cognitive robotics, body schema, forward models, sensorimotor contingencies, grounded cognition, enaction, synthetic methodology, embodied cognition.

1 Introduction

Within the cognitivistic paradigm in cognitive science (e.g., [19,47]), thinking is understood as a result of computation over symbols that represent the world. On the other hand, physical activities, like walking, may be looked at as very low-level, simple and, therefore, uninteresting with regard to the study of cognition. More recently, the view of cognition as symbolic computation has been challenged, and an embodied, action-oriented, dynamic, and developmental view has been proposed instead (e.g., [58,55,39,45,16]). The boundaries between cognitive and non-cognitive phenomena have started to blur and the key influence of the body and the physical interaction with the environment has become accepted. Furthermore, a central role of developmental processes in the emergence

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory,*
Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_15,

of cognition has been asserted. There is growing and increasingly detailed evidence from psychology and neurosciences in support of the *embodied cognitive science* view. However, the premises of the new paradigm—whole brain-body-environment systems rather than isolated subsystems should be studied over extended time periods—pose new challenges to practical empirical research in animals and humans. Here, cognitive developmental robotics as a synthetic approach, i.e. instantiating and studying the phenomena of interest in robots, can serve as a viable tool to verify certain hypotheses and complement the research in psychology and neuroscience [44,3,42].

In this article[1], we will first try to categorize the research in robotics from an (embodied) cognitive science viewpoint (Section 2). Then we will use a quadruped robot and investigate the possibilities of its autonomous development from simple reactive to the first cognitive behaviors: from locomotion to cognition (Section 3). The scenarios are chosen such that they can be successfully mastered only if the robot leaves the "here-and-now" time scale of reactive, stimulus-response, behaviors [59]. In order to do that, the robot needs to extract some regularities from its interaction with the environment and utilize them when selecting the next actions to take. In Section 4, we will then attempt to interpret the case studies from two different viewpoints: a grounded cognition [4] or minimal robust representationalist [9] perspective followed by a non-representationalist or enactive account [58,54]. In the case studies, we explored three concepts that were proposed to explain the development and operation of minimal instances of cognition: body schema (e.g., [29,12]), forward internal models (e.g., [61,11]), and sensorimotor contingencies (SMCs) [39]. Concrete implementations in the robot help us to better understand the meaning of each of them and implications for cognitive development—this analysis will be the topic of Section 5. Then, the implications and limitations of using robots as cognitive science tools will be discussed (Section 6). In particular, we will investigate, whether robots can serve as an essentially paradigm-neutral research tool, or whether their use poses intrinsic limitations—with regard to enactive cognitive science viewpoint, for example. We will close with a conclusion.

2 A Cognitive Classification of AI and Robotics Research

In this section we will strive to sketch a "cognitive landscape" in order to classify some seminal work in Artificial Intelligence (AI) and robotics from the point of view of cognitive science. Obviously, cognition is a very difficult phenomenon and any attempt to "pin it down" in a single diagram is bound to fail. Nevertheless, we believe that it will still be valuable to depict some of the key facets of cognition and their instantiation in AI and robotics in graphical form. To this end, we propose four different axes and hence two different 2D schematics.

[1] Parts of this article are based on the author's PhD thesis [28].

2.1 Offline Reasoning Capability vs. Real-Time Responsiveness

In the first diagram (Fig. 1), the y-axis essentially follows the "grounded cognition" viewpoint of Barsalou [4], and Clark & Grush [9], who used the capability of offline reasoning (or "environmentally decoupled thought") to demarcate cognitive agents (from non-cognitive agents). However, if the role of cognition is to support purposeful and timely action in the real world, the "cognitive" axis needs to be complemented by another dimension, which we have labeled "real-time responsiveness". In other words, cognition should not come at the expense of fast interaction with the environment.

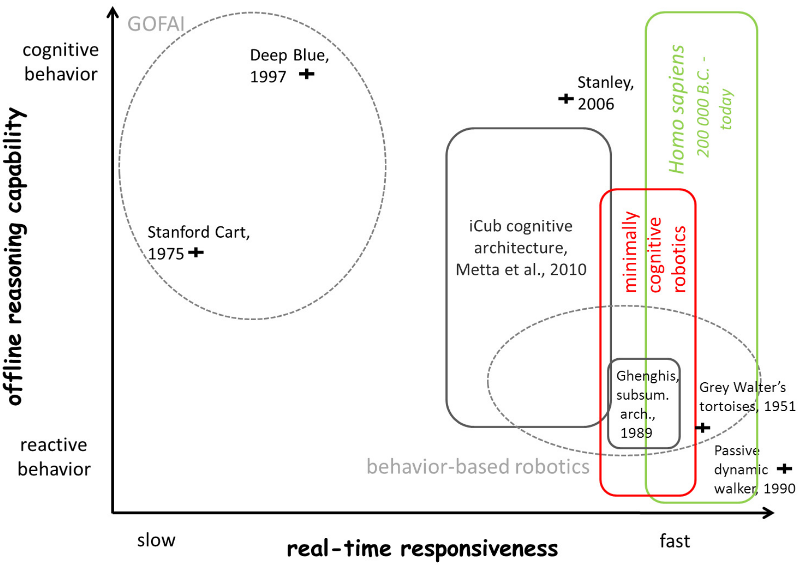

Fig. 1. Cognitive landscape - Offline reasoning capability vs. real-time responsiveness. This figure attempts to classify selected work in AI and robotics from the perspective of cognitive science according to the agents' capability of "running cognition offline" vs. preserving the means to respond immediately. See text for details.

In most machines, their creators decided for a single intersection of these two dimensions—that is the control architecture operates on only one level. These are represented by crosses on the schematics. Passive dynamic walkers [36] can be depicted at the very bottom of the "cognitive axis" (y-axis) as they are passive mechanical machines and are completely coupled to their physical environment. At the same time, their real-time responsiveness is maximum. Examples of reactive agents—creatures capable of simple stimulus-response behavior only—would also occupy the bottom of the "cognitive axis". The tortoises of Grey Walter [60] were composed of direct, analog links connecting sensors and motors; consequently, the real-time responsiveness of the tortoises will be almost as high as that of the passive dynamic walkers. Examples of Good-Old Fashioned Artificial Intelligence (GOFAI), on the other hand, end up on the other end of spectrum: The chess computer, Deep Blue, is definitely capable of offline reasoning (pondering thousands of hypothetical evolutions of the game) and it is subject to soft time constraints—it does not have to respond immediately, yet its total thinking time is limited. From mobile robots, the Stanford Cart [38] was capable of offline reasoning, yet to the extent that it had lost real-time responsiveness (thinking around 15 minutes before every 1 m lurch)[2]. A modern sibling of the Cart, Stanley (DARPA Grand Challenge winner, [57]), is capable of autonomously planning and following a path in an outdoor environment, while preserving very good responsiveness to the current situation on the road.

However, a single degree of "offline reasoning capability"—cognitive, reactive, or only physical—is not sufficient to master a variety of capacities that more complex organisms demonstrate. Therefore, they typically employ a weak hierarchy of different levels ranging from mechanical feedback loops (e.g., [6]) over simple spinal reflexes, which involve direct sensorimotor connections, to more complex and abstracted layers that are present in the brain.

To illustrate this, we have depicted humans ("Homo sapiens") with a large region which ranges all the way from reactive to cognitive behavior on the "cognitive axis". Another example of behavior-based robotics [2], next to Walter's tortoises, is the robot Ghenghis [7]. The so-called subsumption architecture consists of different layers—all of them essentially reactive. Hence, it is also depicted with a small region in the reactive domain rather than a single cross. The cognitive architecture of the iCub humanoid robot as presented in [37] also contains "reactive layers", but at the same time, certain modules reach out of the simple stimulus-response realm to more decoupled processing. Both—Stanley and iCub—are capable of offline reasoning, while preserving real-time responsiveness. Alas, considerable computational resources are required. Finally, "minimally cognitive robotics"—the focus of our interest—would correspond to less "offline reasoning" and more "real-time responsiveness" than Stanley or iCub, building directly on top of the behavior-based robotics school.

2.2 Nature of "Neural Vehicles" and Their Plasticity

In the second diagram (Fig. 2), we propose two additional axes. The y-axis depicts the nature of the internal informational structures that mediate the agent's interaction with the world. They were called "neural vehicles" by Engel [15], avoiding the problematic label of "representation" ("neural" is not be taken literally and is synonymous with internal or belonging to the controller; more details will be provided in Section 4.2). The axis spans the space from no neural vehicles over internal structures operating on the sensorimotor space to symbolic spaces. The x-axis characterizes the degree to which the system has been engineered and remains fixed afterwards or—at the left end of the

[2] It should be noted that this was largely due to the computational power available at that time.

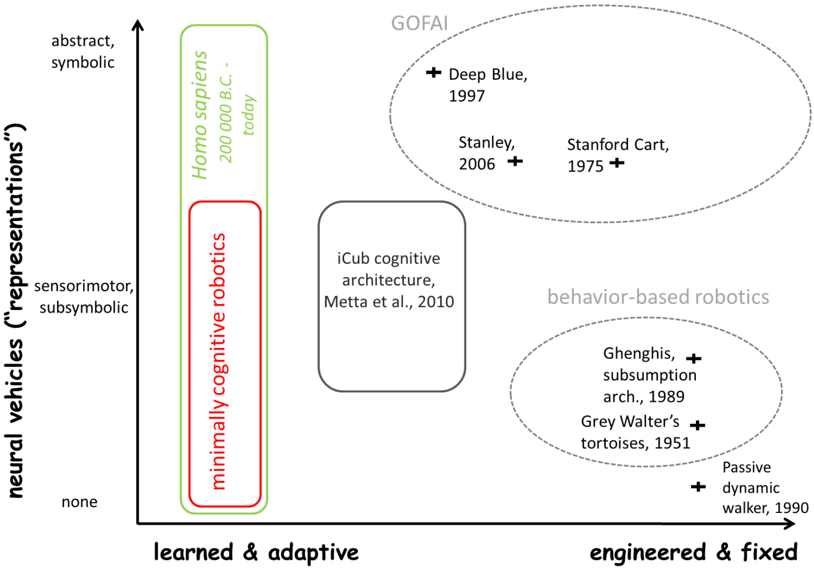

Fig. 2. Cognitive landscape - nature of "representations" vs. their plasticity. This figure attempts to classify selected work in AI and robotics using two additional axes compared to Fig. 1: the character of the control structures vs. their adaptivity. See text for details.

axis—how much was learned without prior knowledge and how adaptive or plastic the neural vehicles are.

Humans are presumably relying on a multitude of mechanisms from direct physical interaction with the world, to simple reflexes and spinal control, to neural vehicles that operate on highly abstracted levels. Hence, they essentially span the whole range of the y-axis. Their location on the x-axis depends on the stance one takes in the "nature vs. nurture" debate. Some capabilities were learned on the evolutionary time scale, which—from the point of view of the individual—represent the "design"/engineering. Yet, the adaptive capacities of humans are extraordinary and support the positioning toward the left end of the axis.

Obviously, there are no neural vehicles, or representations, in the passive dynamic walker. The other instances of behavior-based robotics rely on simple connections of sensors and motors and are thus approaching the "sensorimotor region" from the bottom. On the x-axis, they all have the same value, as they were simply designed by their creators. The GOFAI examples fall on the opposite end of the "representation" axis, since they rely on very abstract or symbolic internal structures. On the "learned vs. engineered" axis, the mobile robots Cart and Stanley display only a limited degree of plasticity (data-driven parameter tuning was applied in [57]). Deep Blue is the most adaptive in this respect, as it learned the position evaluation function itself from thousands of grandmaster games. The iCub's cognitive architecture [37] consists of numerous modules, the core competences being centered around the sensorimotor level. The modules feature different degrees of engineering vs. learning. Finally, the space that is uninhabited by robots so far and that lies at the center of our interest in this article is the "minimally cognitive robotics" region. We will present case studies that attempt to "colonize" this space by learning basic sensorimotor capacities, including integration of information over time and its deployment, starting from minimal prior knowledge.

2.3 Cognitive Developmental Robotics

The iCub humanoid robot and its positioning in our diagrams is a representative example of cognitive developmental robotics, which can be, for example, defined as follows:

> Cognitive developmental robotics (CDR) aims to provide new understanding of how human higher cognitive functions develop by means of a synthetic approach that developmentally constructs cognitive functions. The core idea of CDR is "physical embodiment" that enables information structuring through interactions with the environment, including other agents. The idea is based on the hypothesized developmental model of human cognitive functions from body representation to social behavior. [3]

CDR is thus a subset of developmental robotics in general, which has the same mission, but is not concerned with cognitive phenomena only (however, we have to keep in mind that the boundary between sensorimotor and cognitive phenomena is blurred). A review of developmental robotics is provided by Lungarella et al. [32] or by a special issue of the Infant and Child Development Journal [46]. A review of CDR is provided by Asada et al. [3].

What we have labeled "minimally cognitive robotics" can be seen as a special subset of developmental and cognitive developmental robotics that is specifically concerned with minimal settings where the first instances of offline reasoning—or learning from experience to avoid commitment to the representationalist standpoint—capabilities emerge. In the next section, we provide an overview of experiments that we have performed in a quadruped robot that demonstrate such a developmental pathway.

3 A Developmental Pathway in a Quadruped Robot

In this section, we will present a selection of the results we obtained by instantiating a "cognitive development pathway" in a quadruped robot. The case studies presented feature the key ingredients that are believed to be necessary for cognition to emerge: rich body dynamics and physical interaction with different environments, active generation of multimodal sensory stimulation and learning from this experience over different time scales. In a first study (not reported here, [26]), the robot first learned coordinated movement commands (gaits somewhat resembling those seen in nature – walk, bound, pace etc.) which later formed its motor repertoire. The other case studies deal with the sensorimotor space and the possibility for the robot to extract regularities in it and later exploit this experience in accordance with its goals. More details can be found in [28] and in individual articles reporting the results.

3.1 Specifics of Cognition in a Quadruped Robot

The main platform in our work was the quadruped robot Puppy (Fig. 3). An obvious implication of the embodied cognition stance is that the kind of cognition that will emerge will be highly dependent on the body of the agent, its sensorimotor apparatus and the environment it is interacting with. In our case, the multimodal sensory set together with the nonlinear, partly passive, dynamics of the body can be exploited to extract information about the body itself and the environment. In addition, the absence of distal sensors (camera) forces the robot to use all the modalities by actively probing the environment, which is in accordance with the action-based view on perception and cognition.

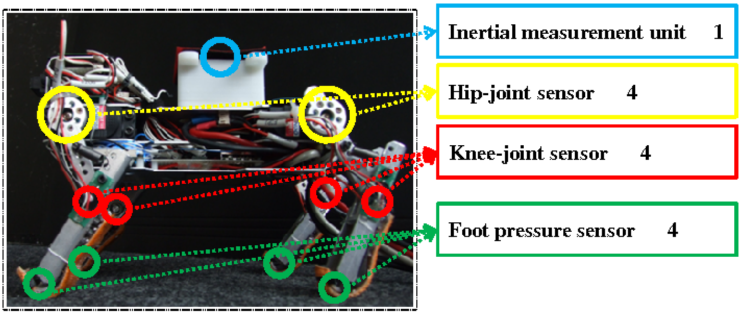

Fig. 3. The quadruped robot Puppy and its sensors. It has four active revolute joints controlled by servomotors ("shoulders" and "hips", in what follows simply hips) and four passive revolute joints at the "elbows" and "knees" (later simply knees)—the passive joints have springs attached across them, making them compliant. There are angular position sensors in all the joints. In addition, there are pressure sensors on the robots feet and an inertial measurement unit (IMU – 3-axis accelerometer, 3-axis gyroscope) on the back. The robot is 20 cm long. **Labeling of channels** (to be used below). The four legs are abbreviated as FL: fore left, FR: fore right, HL: hind left, and HR: hind right. Then, M_{FL}, M_{FR}, M_{HL}, M_{HR} correspond to the four motor channels; H_{FL}, H_{FR}, H_{HL}, H_{HR} denote potentiometers in the hip joints, and K_{FL}, K_{FR}, K_{HL}, K_{HR} in the passive knee joints; P_{FL}, P_{FR}, P_{HL}, P_{HR} are feet pressure sensors, A_X, A_Y, A_Z linear accelerations in three axes, and G_X, G_Y, G_Z are angular velocities. (Figure adapted from [51].)

The locomotion context is particularly suited for understanding minimally cognitive behavior. Whereas "manual cognition", i.e. reaching, grasping and dexterous manipulation, is largely restricted to humans and primates, "locomotor cognition", on the other hand, can be found in much lower animals. For example, path integration was discovered in ants [64]; prediction was demonstrated in motor preparation of prey-catching behavior of a jumping spider [52]; frogs were found to be able to predict whether an aperture could be passed [10]; finally, rats were found covertly comparing alternative paths in a T-maze, thus "planning in simulation" [23] (see [41] for a review). In this work, we will present the robot with similar scenarios: path integration, terrain discrimination and gait selection, and catching another robot.

3.2 Extracting a Body Schema from Raw Sensorimotor Data

In the first study (for details see [51]), we let the robot apply different motor patterns and recorded the corresponding sensory stimulations from its multimodal sensory set comprising primarily tactile and proprioceptive sensors. Then, we systematically analyzed the directed information flows between motors and sensors and showed how the robot could infer a primitive map of its body by extracting the structure of the sensorimotor space that is invariant to changes of the controller: A random set of motor commands proved the most effective in this respect. An information theoretic method that quantifies directed information flows between two variables (sensory and motor time series in our case), transfer entropy, was used. The result is depicted in Fig. 4.

Unlike the majority of work on automatic model acquisition in robotics—reviewed in [24]—, which typically builds on significant prior knowledge and only refines an existing representation using vision, our method is purely data-driven and extracts the regularities intrinsic to the robot's morphology from scratch. Furthermore, the same approach can be used to move from an initial "synesthetic state"—with undifferentiated sensory modalities—to an unsupervised discovery that there are qualitatively different types of sensors.

- **Proprioceptive vs. exteroceptive modality as a graded distinction.** First, we looked specifically at information flows from motor to sensory channels. Those channels that receive strong directed information from the motor signals can be said to be "controllable" by the robot and thus reflecting the state of the the body (under the interpretation "my body is what is under my control"). Hence, these sensors can be said to have "proprioceptive" properties. Exteroceptors, on the other hand, can be defined as sensory channels sensitive to environmental changes.[3] Applying these definitions to the information flows that the agent measured gives a graded distinction of the sensors (see Fig. 5 (left)). Interestingly, only the angular position sensors in the motor-driven hip joints fell clearly into the "proprioceptive" region.

[3] We have systematically varied the environmental conditions—grounds of different friction—and analyzed the data. Please see [51] for the details.

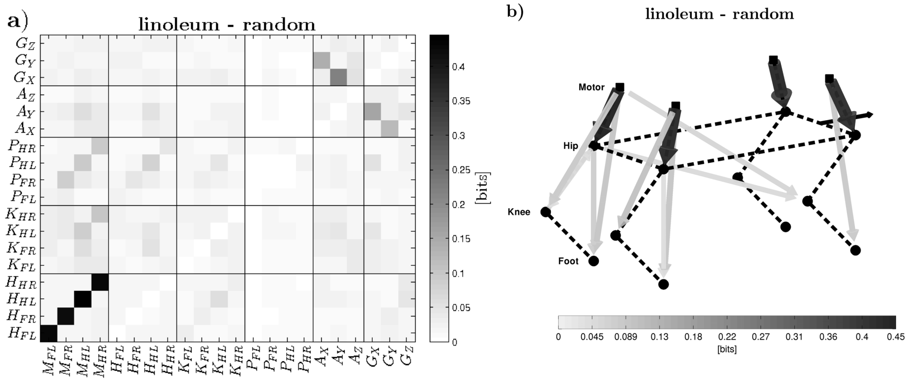

Fig. 4. Transfer entropy TE between all pairs of motor and sensory channels using random motor commands on linoleum ground. (a) Every cell of the matrix corresponds to the information transfer from the signal on the column position to the signal on the row position. Cf. Fig. 3 for the labeling of channels. **(b)** A schematic of the Puppy robot (dashed lines) with overlaid arrows depicting the TE between the individual components. For readability, only the 15 highest values are shown and the accelerometers and gyroscopes were excluded from this visualization. The strength of the information transfer is encoded as thickness of the arrows. The strongest information transfer occurs from the motor signals to their respective hip joint angles ($M_{FL} \rightarrow H_{FL}$, $M_{FR} \rightarrow H_{FR}$, $M_{HL} \rightarrow H_{HL}$, $M_{HR} \rightarrow H_{HR}$). The motors directly drive the respective hip joints and, despite some delay and noise, the hip joints always follow the motor commands, which induces a strong informational relation. The motors further show a smaller influence on the knee angles (especially at the hind legs K_{HL} and K_{HR}) and on the feet pressure sensors, all on the respective leg where the motor is mounted, thus illustrating that body topology was successfully extracted (Figure from [51].)

The other sensors—most of which would be labeled as proprioceptors using a standard "textbook" definition—were found to be more sensitive to the environment.

– **Learning about different sensory modalities.** According to O'Regan and Noe [39], it is the SMCs, i.e. the structure of the rules governing the sensory changes produced by various motor actions, what differentiates modalities. We have applied a similarity measure to the information flows and projected the sensors and motors to a 2D space, creating a sensoritopic map. The resulting map (Fig. 5 (right)) shows a reasonable clustering of angular sensors in active vs. passive joints, pressure sensors, and inertial sensors—reconfirming the SMC hypothesis and demonstrating that no additional knowledge is necessary. The motor modality, which has a different "causal content", is completely separated out on the right of the map.

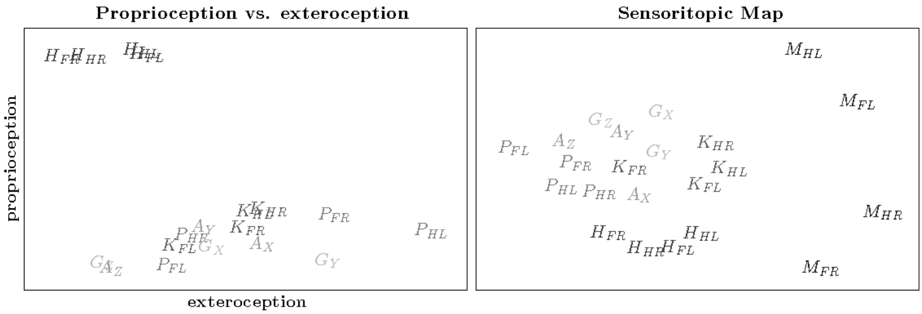

Fig. 5. Sensor spaces in the Puppy robot. (left) Proprioception vs. exteroception. **(right)** A Sensoritopic map. Projection of the sensors and motors into 2D space using multidimensional scaling based on an information flow-based similarity measure. Cf. Fig. 3 for the labeling of channels. (Figure from [51].)

3.3 Learning from Sensorimotor Experience

Whereas the previous case study had an analytical focus—how salient relationships in the sensorimotor space can be extracted—, the next logical step on the "cognitive ladder" is to take the agent's perspective and study how it can integrate its experience and use it to improve behavior. We conducted three studies in this direction, presenting the robot with different tasks that can be successfully mastered only if the robot learns from past sensorimotor experience.

Path Integration Using Self-motion Cues. Humans, other mammals, and also arthropods are reported to be able to perform path integration: estimating the distance traveled without relying on an external reference [17,14,64,65]. Odometers (step integrators) were found to play an important part in this capability. To estimate the length of the step (or stride), the animal seems to require a body representation of some sort ([65] mention: knowledge about intrinsic dynamics of limb segment motion, relationships between gait parameters and body proportions).

In our quadruped robot, we developed one possible solution to the problem: an implicit (data-driven, black-box) model that linearly combines features from multiple sensors from the robot's legs to a stride length estimate [48,49]. Sensory features that correlated most strongly with stride length were selected and a linear regression function that combined them into a stride length estimate was derived, giving rise to a multimodal legged odometer. That is, we showed an example of a procedure that can be employed by an autonomous agent: investigate relationships between a variable of interest and the sensory (or sensorimotor) space, select the signals with the strongest relationships, and work them out into a function. The stride length estimates can then be aggregated over time, giving rise to a measure of distance traveled by the agent—a first example of *integration of information over time* in our agent.

Using Sensorimotor Contingencies for Terrain Discrimination and Adaptive Walking. In this study [27], a record of *past experience in the sensorimotor space was used to inform action selection*: the robot learned to estimate the effects of the application of different gaits in different contexts and used this information to choose the actions that maximize a reward signal (fast and stable locomotion). Eventually, it learned to select an appropriate subset of gaits in different contexts (see [27] for details). No abstraction or hierarchy was used, but a memory of almost raw sensorimotor sequences (compressed into features) allowed the robot to detect familiar contexts and select actions accordingly. Furthermore, we want to highlight two additional outcomes of this study:

– **Perceptual categorization from sensorimotor sequences.** Perceptual categorization can be simplified through embodied interaction with the environment and active generation of sensory stimuli (see e.g., [43]). In our study, when the robot was running on different grounds, only certain, prestructured, stimuli were inevitably induced in the sensory modalities. In addition, the particular action used at every moment—the gait—co-determined what was sensed. We demonstrate this effect by showing the improvement in ground classification when data generated by different gaits are classified separately. Furthermore, we again confirm the hypothesis (put forth in [39] and tested in a simple robot in [35]) that object categorization (the ground being the object here) is improved if longer sensorimotor sequences are considered. The data from both real and simulated robot convincingly demonstrate this (Fig. 6).

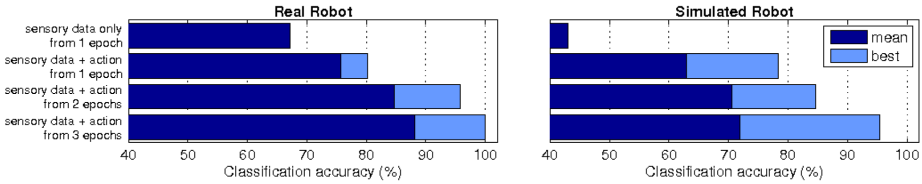

Fig. 6. Comparison of ground classification accuracies when the action context is taken into account to different degrees. The first row corresponds to data from one sensory epoch collapsed across all gaits, i.e. without the action context. Subsequent rows report results where classification was performed separately for each gait and increasingly longer histories were available. "Mean" values represent the mean performance over the individual classifications runs preconditioned on the gait the robot was using; "best" are classification results from the gait that facilitated perception the most. (left) Real robot. (right) Simulated robot. (Figure and caption from [27].)

– **Compression of sensorimotor space through embodiment.** We employed a model presented in [35] and adapted it to our situation. An exhaustive approach to remembering sensorimotor experience was used: the agent did not try to explicitly extract the structure of the sensorimotor space and store it in a compressed form; instead, every new action-observation combination and their history of up to 4 epochs (10 seconds in total) was added to the memory. Although the theoretical dimension of the sensorimotor space was enormous, due to the constraints imposed by the morphology of the robot's mechanical and sensory system, the nature of the interaction with the environment, the action repertoire, and the action selection algorithm, only a small portion of the theoretical state space was visited (2 to 4% of possible states; see [28] for the details). This is in accordance with previous findings on how sensorimotor information is structured through embodiment [33]. That is, the regularities in the sensorimotor space assist the robot in dealing with the curse of dimensionality.

Moving Target Seeking with Forward and Inverse Models. This study [40] constitutes our last step of incremental cognitive development in a mobile robot. We prepared a scenario in which a "hunter" robot needs to catch its conspecific "prey" robot[4]. The scenario was manipulated in order to investigate under which conditions more elaborate planning becomes necessary and what are the best candidates for the implementation. The "hunter" robot was progressively forced by the task-environment to employ less reactive and more cognitive strategies. Finally, it arrived at a multi-step planning architecture: *a "decoupled" forward model, which can be executed independently.* This corresponds to the "cognitive hallmark" proposed by Clark and Grush [9]. The specific points addressed were:

– **Learning a forward model.** A forward model predicting the robot's change in position and orientation was learned through random exploration of the effects of different gaits. An egocentric reference frame was used and no prior knowledge about the platform (such as its kinematics or dynamics) were necessary.
– **Goal state and inverse modeling.** In order to reach the goal state—coming as close as possible to the prey robot—, an inverse model became necessary. That is, given a current position and orientation and a desired one, the output was the best action to take. This was obtained through simple Bayesian inference.
– **Multi-step planning.** We presented the robot with different scenarios: whereas a simple application of the inverse model yielded satisfactory results in some scenarios (hunter and prey in a wall-enclosed arena), in others (open environment) it did not suffice. There, we studied how multistep planning can improve the results. In order to cope with a combinatorial explosion of possibilities, a heuristic best-first-search was implemented.
– **Extending modeling to other agents.** Finally, the utility of explicitly modeling a part of the environment (the "prey" agent) was evaluated and

[4] In this study, simulated Khepera robots with a discrete "gait repertoire" to mimic the situation in the quadruped robot were used. Perception of the hunter's and prey's position were simplified and "GPS" signal available in the simulator was used. Experiments on simulated models of the Puppy robot are under way.

successfully incorporated when it improved the agent's performance. In this way, the agent extended the space of its "cognitive processes" to other agents.

4 Two Sides of the Same Coin? A "Grounded Representation" vs. a Non-representationalist Perspective on the Case Studies

In a nutshell, the case studies presented are concerned with the structure of the sensorimotor space: How it is shaped by an agent's body and dynamic interaction with the environment and how invariant relationships can be extracted and exploited by the agent to improve its behavior. While, on one hand, the performance of the robots in the tasks "speaks for itself", there are still many conceptual questions pending. In particular, should extracting and exploiting past sensorimotor experience be equated with the notions of storage, knowledge, representation, or offline reasoning? In this section, we will attempt something that is rarely undertaken: We will interpret the very same results from two different perspectives: one that posits representations followed by one that rejects them.

4.1 Increasing the Offline Reasoning Capability from Bottom-Up: A Minimally Representationalist Account

The case studies presented lend themselves easily to an interpretation along the lines of "grounded cognition" [4] and "minimal robust representationalism" that was proposed by Clark & Grush [9] in defense of the notion of representation in cognitive science as well as robotics. This view essentially suggests that through processes of internalization and decoupling, cognition can eventually "run in the brain" [9,63].

The "Extracting a body schema" case study (Section 3.2) can be naturally viewed as the robot building a sensorimotor representation of its body—a correspondence between the structure it learns (the representation) and the physical and sensorimotor properties of the robot (what is being represented) can be established. This primitive, sensorimotor structure is extracted solely from the sensory and motor channels and is thus automatically grounded. In the "Path integration" study, the robot builds what could be called a "locomotor body schema", i.e. a model of how much distance it covers every stride. Yet, this "legged odometer" was tuned by using an external reference frame; thus, the "grounding" of this representation—position and orientation in a Cartesian reference frame—is mainly on the side of the observer. The "Terrain discrimination and adaptive walking study" lends itself to a representational interpretation too. The sensorimotor histories that are stored in the associative memory ("the brain" of the agent) can be looked at as knowledge or representation of the robot's previous interactions with the world. They can also be used to classify the environments, replayed or iterated forward to get predictions and inform action selection—in accordance with the "offline cognition" notion. Finally, the

"Moving target seeking" study serves as a perfect example of a bottom-up development of an internal simulator/emulator [9,22]: the robot learns a forward and inverse model of the outcome of applying different motor commands. Later—in order to succeed in the task of catching its conspecific—it also learns a model of the "prey's" behavior and applies a multi-step planning algorithm. This demonstrates an increasing degree of offline reasoning and matches with the evolutionarily plausible path how internal representation (in the form of emulator circuitry) could possibly get its foot in the door of real-world, real-time cognition [9].

4.2 Enactivism and "Cognition-is-not-in-the-Brain" Viewpoint

Interestingly, we can try to embrace the very same case studies into a more radical school of thought that rejects the neurocentric perspective on cognition altogether. A unique perspective on cognition has been offered by the community that has grown around the work of Francisco Varela (e.g., [58]). The proponents of the *enactive framework* reject the idea that "cognition often proceeds independently of the body"[4]. For the "enactivists", cognition is not only shaped by the body and its action possibilities, but *cognition is action*—embodied action, a form of practice itself [58]. In this view, cognition is not about world-mirroring through representations, but "world-making" and sense-making. The interested reader is referred to the abundant literature (e.g., the recent collection of papers in [54], reviewed in [20]).

We will borrow useful terms from Engel et al. [15,16] who provide a review of a turn toward action in cognitive science and propose the term *dynamic directive* to "denote the action-related role of large-scale dynamic interaction patterns that emerge in a cognitive system. On this account, directives can be defined as dispositions for action embodied in dynamic activity patterns." Importantly, the directives are not equal to states in the brain (and thus are not equal to action-oriented representations; see also [30] for a detailed philosophical account of this distinction), but refer to dynamics of the whole—or relevant parts of a—brain-body-environment system. At the same time, it may be convenient to invoke a term for the "traces" of the directives in the brain: These are the *neural vehicles* of the directives [15].

Some of the results we have presented are not compatible with this viewpoint. For example, the path integration case study was devoted to the learning of a position and orientation estimation module trained by an external reference. Moreover, in this particular study, no action selection was performed based on the path integration results. Thus, this task had little significance or meaning for the robot. In summary, the focus was on a veridical representation of the position of the robot in the environment—an emphasis that is incompatible with the formulations that belong to enactive cognitive science.

Let us look at the "adaptive walking" and the "moving target seeking" studies. There, the robot had to optimize its behavior on a task—fast and stable walking in the former case, catching another robot in the latter. To this end, different control architectures that could assist the robot in the task were explored. In the predator-prey scenario, the "neural vehicles" were data-driven and learned *ab initio*, but the structure of the model (a simple Bayesian network), the variables of interest (distances and angles), and the goal (catching the prey) came from the designers. The "world-making" of the robot has thus been relatively strongly constrained and imposed on the robot from the outside. Finally, the "adaptive walking" study, where a model of sensorimotor contingencies is employed, could probably be most in line with an enactive viewpoint. The robot simply records past sensorimotor experiences (the gait used and all the sensory channels compressed into features) together with the values of the reward function and uses this information to inform future decision-making: selecting the gait that is most likely to succeed on a given ground. The "neural vehicle" thus contains raw sensorimotor "footprints" of the robot's interaction with the environments and uses them for action guidance. The individual terrains are nowhere explicitly coded in the neural substrate—they are implicitly recognized by selecting appropriate actions. Yet, the reward function was again defined from the outside and the "sensorimotor look-up table" that is driving the behavior at discrete time steps is perhaps still too decoupled from the dynamics of the body and environment when compared to the—alas much simpler—dynamical accounts of active categorical perception [5,8][5].

5 Body Schema, Forward Models, and Sensorimotor Contingencies: On Their Overlap, Definition, and Degree of Representational Nature

We have set out to investigate bottom-up development of minimally cognitive abilities. On this path, we have repeatedly encountered three concepts: body schema, forward internal models, and sensorimotor contingencies (SMCs). We have explored them in different disguises in our robotic case studies. What can we now say regarding their nature, utility and compatibility with different cognitive science paradigms?

As reported by Rochat [50], infants spend substantial time in their early months observing and touching themselves. Through this process of babbling, intermodal redundancies, temporal contingencies and spatial congruences are picked up. This basically encompasses all the low-level relationships that an agent can learn during its early development. However, this space is too large. Therefore, in order to bootstrap its development, an agent needs to focus on some subspaces of the sensory-motor-time space and choose an appropriate way of modifying its internal dynamics in accordance with these regularities and its goals. The three aforementioned concepts qualify as suitable candidates in this regard.

[5] Note that the perceptual categorization we performed in Section 3.3 used solely the sensorimotor memory, i.e. the neural structure. Beer or Buhrmann et al. [5,8], on the other hand, show that in their examples, this is not possible.

5.1 The "Minimally Cognitive Concepts" in the Case Studies

– **Body schema.** As we have argued, the body has a key influence on the agent's behavior as well as on the information that enters its brain/controller (see [25] for a collection of examples illustrating this). Therefore, it can bring advantage to the agent if it can pick up the regularities that are induced by its body. The concepts of body schema and body image are used in this context. However, at the moment, they serve more as "umbrella concepts" for a multitude of body representations that animals and humans develop and use (cf. e.g. [12]). The synthetic approach allowed us to explore these concepts in more concrete terms. In the "Extracting a body schema" study (Section 3.2), we investigated two possibilities for the formation of a primitive body representation in a robot. First, we studied the structure of the sensorimotor space that is invariant to changes in the motor commands and the environment—that is, body as the invariant structure in sensorimotor space. Second, we studied which sensory channels were strongly affected by the motor signals. This provides an alternative view: the agent's body is what it can control. Both viewpoints can have merits for the agent: the former one could be used for self-diagnosis (if the invariant structure changes, this can be attributed to changes in the body), the latter one can be used to bootstrap development—learning the first behaviors. Yet, this is just the very beginning and subsequent development needs to be demonstrated. A more narrow type of body schema or image devoted specifically to estimating the robot's stride length was developed in the "Path integration" case study.
– **Forward model.** Forward model is another type of mapping that can be useful to the agent. It can be used to predict the next sensory state (given the current state and a motor command) or—if chained or iterated—even to simulate whole sensorimotor loops covertly. It is concretely defined[6] and can be instantiated at any abstraction level (i.e., not only for low-level motor control, where the existence of forward internal models is subject to a heated debate – cf. for example [11] vs. [18]). We have explicitly employed probabilistic forward and inverse models in the "Moving target seeking" study. The architecture used in the "terrain discrimination and adaptive walking" study that is using conditional probability distributions [35] also encompasses forward model functionality.
– **Sensorimotor contingencies.** Sensorimotor contingencies (SMCs) were originally presented in the influential article by O'Regan & Noe [39] as the structure of the rules governing sensory changes produced by various motor actions. Similarly to a body schema, this notion is still not articulated concretely enough to allow for an implementation in a robot. For example, is a forward model an instance of an SMC? Also, what is the "site" where the SMCs reside—are they stored in the brain? Very recently,

[6] The forward model is classically thought of as a function, $f(s_t, m) = s_{t+1}$, which maps a sensory state and a motor command to a next sensory state (where the states can be multidimensional).

Buhrmann et al. [8] have addressed these questions and proposed a dynamical systems account of SMCs, distinguishing between *Sensorimotor (SM) environment*, *SM habitat*, *SM coordination*, and *SM strategy*. The SM environment is the relation between motor actions and changes in sensory states, independently of the agent's internal (neural) dynamics. Interestingly, this definition closely resembles the forward model that we have encountered before.[7] The other notions—from SM habitat to SM strategies—add internal ("brain") dynamics to the picture. SM habitat refers to trajectories in the sensorimotor state space, but under certain conditions on the internal dynamics that is responsible for generating the motor commands. These are thus not random anymore and may depend on previous sensory states as well—an example of closed-loop control. SM coordination then further reduces the set of possible SM trajectories to those "that occur reliably and contribute functionally to a task". For example, specific patterns of pressing an object in order to assess its hardness would be SM coordination patterns serving object discrimination. Finally, SM strategies take, in addition, a normative framework ("reward" or "value" for the agent) into account.

Taking advantage of this operationalization of the SMC concept, in what disguises can we find SMCs in our case studies? In the study described in Section 3.2[8] random motor commands were applied (hence there was random or no neural dynamics) and the relationships between motor and sensory variables were studied, closely resembling the notion of SM environment.[9] Then, we also studied the relationships in the sensorimotor space when the robot was running with certain coordinated movement patterns: gaits. These were obtained by optimizing the robot's performance for speed or for turning [26] and thus correspond to patterns that are functionally relevant for the robot and even carry a normative aspect. Thus, our findings about the sensorimotor space using the gaits (results shown in [51]) can be interpreted as studying the SM coordination or even SM strategy of the quadruped robot. In the "adaptive walking" study, a similar repertoire of coordinated

[7] The functional form was provided in the previous footnote. However, if the sensory state does not fully define the state of the system—which is likely given that the internal neural as well as environmental variables are ignored—it is easy to imagine that this mapping will not be right-unique and thus, mathematically speaking, seize to be a function. The SM environment is an even more general relation, a superset of multiple forward models. In discrete terms, it would have the form $R(m, s_{t+1} - s_t)$.

[8] Note that the Section's name is "Extracting a body schema from raw sensorimotor data", illustrating the confusion of terms.

[9] The particular details differ though. First, due to the dimensionality of the sensorimotor space, we studied relationships between pairs of variables only. On the other hand, as opposed to SM environment, we included sensory-sensory pairs as well. In addition, we applied a particular information theoretic measure, transfer entropy, which allowed us to assess the amount of directed information transfer between individual variables. In this way, information was compressed and salient relationships could be discovered, but at the same time, it did not contain all the information present in the original data.

gaits was used. While exercising these in different environments, the robot was taking a record of all the combinations of sensory and motor variables (discretized and compressed into features over 2 second intervals). A reward associated with every sensorimotor state was stored and later used to inform action selection. Thus, the items in the associative memory constitute discrete slices that witness and at the same time influence the robot's SM strategies. Buhrmann et al. [8] also highlight how the space of possible sensorimotor trajectories—in the original sensorimotor space—is narrowed down as one goes from SM environment to SM coordination. In our example, we quantified the overall compression of a theoretical full sensorimotor state space as a result of embodiment and the action selection (internal dynamics) in Section 3.3.

5.2 Clarification

Let us now try to directly compare the "cognitive concepts" that we discussed above in terms of their mathematical formulation, representational nature, and site—where they are located. With the help of the analysis that follows, we will fill up Table 1.

- **Mathematical formulation.** As we have argued, a body schema is a very loosely defined notion and to talk about a mathematical formulation is out of question. A forward model, on the other hand, can be defined precisely as a function. SMCs were also defined rather loosely, but acquired a concrete articulation in dynamical systems terms in [8].
- **Representational nature.** The term "body schema" is usually equated with a body representation. It thus seems to imply a "representationalist" view of the mind. Alsmith & de Vignemont discuss this theme in detail in [1]. A forward model is simply a function on motor and sensory variables, which is per se neutral with respect to the "representationalist dispute". Of course, representational nature can be ascribed to it if one posits that this mapping is stored in the brain and "stands in" for some extraneural states of affairs, as done by Clark & Grush [9], for example. The position of SMCT[10] (in its original formulation [39]) on representations was not clear—for sure it was detailed, pictorial representations, "mirrors of the world states", that SMCT was arguing against. Buhrmann et al. [8] in their definition and treatment argue clearly against a representationalist interpretation and show that the SMCs—as trajectories in the sensorimotor space—are emergent from the dynamics of the body, brain, and environment (similarly to the dynamic directives proposed in [16]).
- **Site.** A body schema is usually thought to reside in the brain—even if in a highly distributed manner, encompassing for example area SI, area 5 in the parietal lobe, and premotor cortex [21]. The existence of forward models in the brain is also supported by extensive literature, in particular on the

[10] Sensorimotor Contingency Theory.

cerebellum (e.g., [31,11]). SMCs are a result of the joint dynamics of the brain, body, and environment; an analysis of the simple agent in [8] reveals that “there is nothing in the internal dynamics of the agent’s “brain” that represents the SMCs that are being enacted or the non-actualized sensorimotor regularities that still have a dynamical influence.” Yet, some neural vehicles that support the SMCs on the part of the brain seem inevitable and are expected in various brain areas—visual SMCs are discussed in [39]; Engel et al. [16] discuss the role of premotor circuits, for example.

Table 1. Properties of different minimally cognitive concepts

	Body schema	Forward model	SMCs (according to [8])
Mathematical description	N.A.	Function	Trajectory in S-M space
Representational nature	Yes	Neutral	No
Site	Brain	Brain	Brain-body-environment

6 Robots as Cognitive Science Tools: Are There Intrinsic Limitations?

The methodology adopted in this work was a synthetic one [44]. That is, we built and then studied the behavior of artifacts. As can be seen in Fig. 7, the area spanned by synthetic sciences can be further subdivided into (1) the intersection with empirical sciences—synthetic modeling, (2) the middle area concerned with general principles, (3) the intersection with the application domain in the form of prototypes of new technology.

The scenarios presented here (Section 3) were inspired by skills that were observed in lower animals and serve as instances of the simplest behaviors that we would consider cognitive. Yet, do the case studies presented in this work qualify as synthetic modeling, i.e. as models of biological cognition too? Given that we do not treat cognition as an exclusively biological phenomenon, this possibility is open. However, the parallel between biological cognitive agents and the artificial ones remained on an abstract level—we did not relate directly to any empirical data from the animal kingdom. Along the lines of the critical account in [62], one could argue that this is an example of the “animat” approach to modeling cognition and that more direct parallels to concrete instances of cognitive phenomena in biology are desirable. Several proposals in this direction are put forth in [42]. This would be one possible direction of future work (sketched in Section 8.3.3 in [28]).

6.1 The Difficulty of Modeling without Representations

Interestingly, we find that the ease of the synthetic modeling endeavor will depend on the cognitive science paradigm that one is following. Under cognitivism / GOFAI, the body and interaction with environment was of marginal

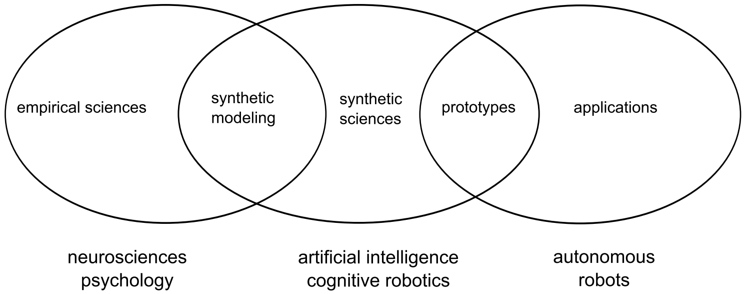

Fig. 7. Overview of approaches to the study of cognition. The figure and caption are adapted from [44] to the study of cognition rather than intelligence. On the left, we have the empirical sciences like neurosciences and psychology that mostly follow an analytic approach. In the center, we have the synthetic ones like AI and cognitive robotics which can either model natural agents (this is called synthetic modeling—the intersection with empirical sciences) or alternatively can simply explore issues in the study of cognition without necessarily being concerned with natural systems. This activity may give rise to prototypes and eventually to full industrial applications, such as autonomous robots.

importance, so the robots were not necessary in the first place. Moreover, the representations—models of the world—were often symbolic and directly corresponded to objects in the world (in their designers' eyes). The quality and functioning of the cognitive layer was thus easy to assess. This was obstructed slightly under connectionism, as the models became less transparent due to their subsymbolic nature. Embodied cognitive science then brought about the necessity for considering whole brain-body-environment systems. However, even within embodied cognitive science, the different viewpoints that we have outlined in Section 4 impact on the research methodology. First of all, the move away from veridical to action-oriented or context-dependent representations means that the quality of the internal control structures cannot be assessed by a direct comparison with some objects in the world anymore. The viewpoints that reject representations altogether go even further in this and imply that no answers will be found in the control structure alone [5,8]. This is analogous to the situation in neurosciences where Engel et al. [16] propose to replace techniques studying neural responses to passive stimuli by studying subjects actively interacting with their surroundings, which brings about many practical difficulties.

6.2 Enactive Robots Subject to Precarious Conditions?

The enactive viewpoint can be taken even further: Di Paolo [13] points out that in order to fully understand cognition in its entirety, embedding the agent in a closed-loop sensorimotor interaction with the environment is necessary, yet may not be sufficient in order to induce important properties of biological agents such as intentional agency. In other words, one should not only study instances of individual closed sensorimotor loops as models of analogous loops in biological agents—that would be the recommendation of Webb [62]—but one should also try to endow the models (robots in this case) with similar properties and constraints that biological organisms are facing. In particular, it has been argued that life and cognition are tightly interconnected [34,56] and a particular organization of living systems—which can be characterized by autopoiesis [34] or metabolism for example—is crucial for the agent to truly acquire the meaning in its interaction with the world. While these requirements are very hard to satisfy with the artificial systems of today, Di Paolo [13] proposes a way out: robots need not metabolize, but they should be subject to precarious conditions. That is, the success of a particular instantiation of sensorimotor loops or neural vehicles in the agent is to be measured against some viability criterion that is intrinsic to the organization of the agent. The control structure may evolve over time, but the viability constraint needs to be satisfied, otherwise the agent "dies". The unfortunate implication, however, is that research along these lines will hardly fit into the full synthetic methodology scheme (Fig. 7) anymore, since machines whose functioning is not deducible from their control structure and that cannot be given tasks will not easily find their way to application scenarios in industry. On the other hand, this approach may give rise to truly autonomous robots.

7 Conclusion

We focused on autonomous cognitive development and engaged robots in a number of scenarios that can be seen as a developmental pathway from reactive to minimally cognitive behavior. We have experimented with different control architectures and assessed their performance in different tasks. We have also analyzed the nature of these control architectures from the point of view of different cognitive science paradigms. We found that our case studies lend themselves easily to interpretations along the lines of "grounded representation" and internal simulation/emulation theories [4,9,22]. On the other hand, if one looks into the details, they are much less compatible with the non- (or anti-) representational or enactive perspectives [58,54].

The minimally cognitive "building blocks" or notions were also subject to investigations in our case studies. Our results and analysis contributed to a conceptual clarification here. Interestingly, only a forward model seems to be a useful building block that can be deployed in the control structures of robots and serve different purposes—a kind of useful "brain motif" [53] perhaps. A body schema is at the moment an "umbrella term" for a multitude of body representations that can be used for action. This notion is, however, far from a formulation that could be "deployed" in a control architecture. Similarly, SMCs do not constitute a building block either; instead, at the moment, they are rather a descriptive concept, which may prove useful in the analysis of natural and artificial cognitive systems.

Finally, we have evaluated the potential of robots as modeling tools for cognitive science and the implications of this way of modeling regarding the choice of cognitive science paradigm. Adopting an embodied, yet representation-based view, is a convenient choice that creates bridges between the research in psychology, neuroscience and robotics (as elaborated recently under the "grounded cognition" umbrella in [42]). In line with the synthetic approach and a functionalist stance, a particular cognitive architecture may serve as a model of certain parts of the brain and at the same time provide an interesting tool for autonomous robotics, for example. Still, it remains to be shown if human-like levels of complexity can be attained. On the other hand, truly enactive robots seem to be much harder to realize. Models that are compatible with this view are to date of minimal complexity and bear no application potential. From a designer's perspective, achieving an appropriate "shaping of dynamical tendencies that channel appropriate actions on the basis of past experience and in accordance with goals" [8] seems to be much harder than adopting the representationalist stance and tuning a world model of one form or another. Therefore, synthetic enactive approaches in robotics still need to demonstrate their scalability and potential.

Acknowledgement. This work was supported by the project 'Extending sensorimotor contingencies to cognition - eSMCs', FP7-ICT-270212, `esmcs.eu`, and by the Swiss National Science Foundation Fellowship PBZHP2-147259. I would like to thank the collaborators involved in the case studies reviewed in this work, Rolf Pfeifer for continuous support and guidance during this research, and to Giovanni Pezzulo, Alexander Maye, and the two anonymous reviewers for helpful and constructive comments on the manuscript.

References

1. Alsmith, A.J.T., De Vignemont, F.: Embodying the mind and representing the body. Review of Philosophy and Psychology 3(1), 1–13 (2012)
2. Arkin, R.C.: Behavior-based Robotics. MIT Press (1998)
3. Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M., Yoshida, C.: Cognitive developmental robotics: a survey. IEEE Transactions on Autonomous Mental Development 1(1), 12–34 (2009)
4. Barsalou, L.: Grounded cognition. Annual Review of Psychology 59, 617–645 (2008)
5. Beer, R.D.: The dynamics of active categorical perception in an evolved model agent. Adaptive Behavior 11, 209–243 (2003)
6. Blickhan, R., Seyfarth, A., Geyer, H., Grimmer, S., Wagner, H., Guenther, M.: Intelligence by mechanics. Phil. Trans. R. Soc. Lond. A 365, 199–220 (2007)
7. Brooks, R.A.: A robot that walks: Emergent behaviors from a carefully evolved network. Neural Computation 1, 153–162 (1989)
8. Buhrmann, T., Di Paolo, E., Barandiaran, X.: A dynamical systems account of sensorimotor contingencies. Front. Psychol. 4(285) (2013)
9. Clark, A., Grush, R.: Towards cognitive robotics. Adaptive Behaviour 7(1), 5–16 (1999)
10. Collett, T.: Do toads plan routes? a study of the detour behavior of bufo viridis. J. Comparative Physiology 146, 261–271 (1982)
11. Davidson, P.R., Wolpert, D.M.: Widespread access to predictive models in the motor system: a short review. Journal of Neural Engineering 2, 313–319 (2005)

12. de Vignemont, F.: Body schema and body image - pros and cons. Neuropsychologia 48(3), 669–680 (2010)
13. Di Paolo, E.: Robotics inspired in the organism. Intellectica 53-54, 129–162 (2010)
14. Durgin, F.H., Akagi, M., Gallistel, C.R., Haiken, W.: The precision of locometry in humans. Exp. Brain Res. 193, 429–436 (2009)
15. Engel, A.K.: Directive minds: how dynamics shapes cognition. In: Enaction: Towards a New Paradigm for Cognitive Science, pp. 219–243. MIT Press, Cambridge (2011)
16. Engel, A.K., Maye, A., Kurthen, M., König, P.: Where's the action? the pragmatic turn in cognitive science. Trends in Cognitive Sciences 17(5), 202–209 (2013)
17. Etienne, A.S., Jeffery, K.J.: Path integration in mammals. Hippocampus 14, 180–192 (2004)
18. Feldman, A.: New insights into actionperception coupling. Experimental Brain Research 194, 39–58 (2009)
19. Fodor, J.A.: The language of thought. Harvard University Press, Cambridge (1975)
20. Froese, T.: From adaptive behavior to human cognition: a review of enaction. Adaptive Behavior (2012)
21. Graziano, M.S.A., Botvinick, M.M.: How the brain represents the body: insights from neurophysiology and psychology. In: Common Mechanisms in Perception and Action: Attention and Performance, vol. XIX, pp. 136–157. Oxford Univ. Press (2002)
22. Grush, R.: The emulation theory of representation - motor control, imagery, and perception. Behavioral and Brain Sciences 27, 377–442 (2004)
23. Hesslow, G.: Conscious thought as simulation of behaviour and perception. Trends in Cognitive Sciences 6, 242–247 (2002)
24. Hoffmann, M., Marques, H., Hernandez Arieta, A., Sumioka, H., Lungarella, M., Pfeifer, R.: Body schema in robotics: a review. IEEE Trans. Auton. Mental Develop. 2(4), 304–324 (2010)
25. Hoffmann, M., Pfeifer, R.: The implications of embodiment for behavior and cognition: animal and robotic case studies. In: The Implications of Embodiment: Cognition and Communication, pp. 31–58. Imprint Academic, Exeter (2011)
26. Hoffmann, M., Schmidt, N., Nakajima, K., Iida, F., Pfeifer, R.: Perception, motor learning, and speed adaptation exploiting body dynamics: case studies in a quadruped robot. In: Proc. Int. Symposium on Adaptive Motion in Animals and Machines, AMAM (2011)
27. Hoffmann, M., Schmidt, N., Pfeifer, R., Engel, A.K., Maye, A.: Using sensorimotor contingencies for terrain discrimination and adaptive walking behavior in the quadruped robot puppy. In: Ziemke, T., Balkenius, C., Hallam, J. (eds.) SAB 2012. LNCS (LNAI), vol. 7426, pp. 54–64. Springer, Heidelberg (2012)
28. Hoffmann, M.: From locomotion to cognition. PhD thesis, University of Zurich (2012)
29. Holmes, N.P., Spence, C.: The body schema and the multisensory representation(s) of peripersonal space. Cogn. Process. 5(2), 94–105 (2004)
30. Hutto, D.: Exorcising action oriented representations: ridding cognitive science of its nazgûl. Adaptive Behavior (2013)
31. Kawato, M.: Internal models for motor control and trajectory planning. Current Opinion in Neurobiology 9, 718–727 (1999)
32. Lungarella, M., Metta, G., Pfeifer, R., Sandini, G.: Developmental robotics: a survey. Connection Science 15(4), 151–190 (2004)
33. Lungarella, M., Sporns, O.: Mapping information flow in sensorimotor networks. PLoS Comput. Biol. 2, 1301–1312 (2006)

34. Maturana, H., Varela, F.: Autopoiesis and cognition: the realization of the living. D. Reidel Publishing, Dordrecht (1980)
35. Maye, A., Engel, A.K.: A discrete computational model of sensorimotor contingencies for object perception and control of behavior. In: Proc. IEEE Int. Conf. Robotics and Automation (ICRA), pp. 3810–3815 (2011)
36. McGeer, T.: Passive dynamic walking. The International Journal of Robotics Research 9(2), 62–82 (1990)
37. Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L., von Hofsten, C., Rosander, K., Lopes, M., Santos-Victor, J., Bernardino, A., Montesano, L.: The icub humanoid robot: An open-systems platform for research in cognitive development. Neural Networks 23(8-9), 1125–1134 (2010)
38. Moravec, H.P.: The stanford cart and the cmu rover. Proceedings of the IEEE 71(7), 872–884 (1983)
39. O'Regan, J.K., Noe, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, 939–1031 (2001)
40. Oses, N., Hoffmann, M., Koene, R.A.: Embodied moving-target seeking with prediction and planning. In: Corchado, E., Graña Romay, M., Manhaes Savio, A. (eds.) HAIS 2010, Part II. LNCS (LNAI), vol. 6077, pp. 478–485. Springer, Heidelberg (2010)
41. Pezzulo, G.: Anticipation and Future-Oriented Capabilities in Natural and Artificial Cognition. In: Lungarella, M., Iida, F., Bongard, J.C., Pfeifer, R. (eds.) 50 Years of Aritficial Intelligence. LNCS (LNAI), vol. 4850, pp. 257–270. Springer, Heidelberg (2007)
42. Pezzulo, G., Barsalou, L.W., Cangelosi, A., Fischer, M.H., McRae, K., Spivey, M.J.: The mechanics of embodiment: a dialog on embodiment and computational modeling. Frontiers in Psychology 2 (2011)
43. Pfeifer, R., Scheier, C.: Sensory-motor coordination: The metaphor and beyond. Robotics and Autonomous Systems 20, 157–178 (1997)
44. Pfeifer, R., Scheier, C.: Understanding intelligence. MIT Press, Cambridge (2001)
45. Pfeifer, R., Bongard, J.C.: How the body shapes the way we think: a new view of intelligence. MIT Press, Cambridge (2007), `http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11003`
46. Prince, C.G. (ed.): J. Infant and Child Development: Special issue: Developmental robotics. Wiley Interscience (2008)
47. Pylyshyn, Z.: Computation and cognition: Toward a foundation for cognitive science. MIT Press, Cambridge (1984)
48. Reinstein, M., Hoffmann, M.: Dead reckoning in a dynamic quadruped robot: Inertial navigation system aided by a legged odometer. In: IEEE Int. Conf. Robotics and Automation (ICRA), pp. 617–624 (2011)
49. Reinstein, M., Hoffmann, M.: Dead reckoning in a dynamic quadruped robot based on multimodal proprioceptive information. IEEE Transactions on Robotics 29(2), 563–571 (2013)
50. Rochat, P.: Self-perception and action in infancy. Exp. Brain Res. 123, 102–109 (1998)
51. Schmidt, N., Hoffmann, M., Nakajima, K., Pfeifer, R.: Bootstrapping perception using information theory: Case studies in a quadruped robot running on different grounds. Advances in Complex Systems J. 16(2-3) (2013)
52. Schomaker, L.: Anticipation in cybernetic systems: a case against mindless anti-representationalism. In: Proc. Int. Conf. Systems, Man and Cybernetics, vol. 2, pp. 2037–2045 (2004)

53. Sporns, O., Kötter, R.: Motifs in brain networks. PLoS Biology 2(11), e369 (2004)
54. Stewart, J., Gapenne, O., Di Paolo, E. (eds.): Enaction: Toward a new paradigm for cognitive science. MIT Press, Cambridge (2010)
55. Thelen, E., Smith, L.B.: A Dynamic systems approach to the development of cognition and action. MIT Press (1994)
56. Thompson, E.: Mind in life: Biology, phenomenology, and the sciences of mind. MIT Press, Cambridge (2007)
57. Thrun, S., et al.: Stanley: The robot that won the darpa grand challenge: Research articles. J. Robot. Syst. 23, 661–692 (2006)
58. Varela, F., Thompson, E., Rosch, E.: The embodied mind: cognitive science and human experience. MIT Press, Cambridge (1991)
59. Vernon, D., Metta, G., Sandini, G.: Embodiment in Cognitive Systems: on the Mutual Dependence of Cognition and Robotics. In: Embodied Cognitive Systems. Institution of Engineering and Technology (IET), UK (2010)
60. Walter, G.W.: The living brain. Norton & Co., New York (1953)
61. Webb, B.: Neural mechanisms for prediction: do insects have forward models? Trends in Neurosciences 27, 278–282 (2004)
62. Webb, B.: Animals versus animats: or why not model the real iguana? Adaptive Behavior 17, 269–286 (2009)
63. Wilson, M.: Six views of embodied cognition. Psychonomic Bulletin & Review 9(4), 625–636 (2002)
64. Wittlinger, M., Wehner, R., Wolf, H.: The ant odometer: stepping on stilts and stumps. Science 312(5782), 1965–1967 (2006)
65. Yvanenko, Y.P., Dominici, N., Daprati, E., Nico, D., Cappellini, G., Lacquaniti, F.: Locomotor body scheme. Human Movement Science 30(2), 341–351 (2011)

Human Language and Sensorimotor Contingency

Stephen J. Cowley

Centre for Human Interactivity, Department of Language and Communication,
University of Southern Denmark
"L'enfer, c'est les autres" J.P. Sartre (1945)

1 Introduction

Using the monumental work of figures like Bloomfield (1933), Harris (1951) and Chomsky (1965), scientific linguistics has often been centrally concerned with verbal patterns. Yet such views have even older roots. Ever since writing systems first arose in Sumeria, what people do as they talk has gradually been standardised in ways that eventually gave rise to the electronic use of artificial codes. First, iconic signs became ideographic or alphabetic. Then, in Europe, mediaeval scribes introduced spaces between units (and written words) and, in the aftermath of new technologies, there came dictionaries, grammars, printing and, yesterday, computers. As a result, written language bias (Linell, 2005) has dominated philosophy, linguistics and classic cognitive science. Languages are seen as verbal systems whose words and rules are, in some sense, *separate* from people. Even talk is often modelled around transcriptions that invite comparison with ways of construing verbal sequences. This can be highly misleading. In fact, while ideographic and alphabetic symbols are unsponsored, speech and hearing *are* coordinated human activity. Language is intrinsic to action and thus partly constitutive of experience: for Sartre (1945), self must be mirrored by others. Although this is a commonplace, many explanatory models reduce doing things with language to how linguistic forms are 'used' and/or 'represented'. In what follows, I adopt the distributed perspective (see, Cowley, 2011b), to offer an alternative.

Language can be seen as a cultural and historical extension of how we exploit our embodiment. On this view, verbal patterns are a small part of language. As argued by Heidegger (1971), Wittgenstein (1958; 1980), Merleau-Ponty (1996) and Gibson (1966; 1979), language is inseparable from action and perception. In cognitive science the view can be traced to Maturana (1978) and, in linguistics, to Lakoff and Johnson (1980). In ecological psychology, speech perception was shown to track phonetic gesture (Browman and Goldstein, 1986) and, elsewhere, language was found to be inseparable from visible gesture (McNeill, 1992). Today, weight falls on coordination, joint activity and acts of meaning (Raczaszek-Leonardi & Kelso 2008; Fusaroli & Tylén, 2012); language needs, not a mental lexicon, but rich phonetic memory (Port & Leary, 2005). Even when alone, the use of vocal and other gestures

J.M. Bishop and A.O. Martin (eds.), *Contemporary Sensorimotor Theory*,

Studies in Applied Philosophy, Epistemology and Rational Ethics 15,
DOI: 10.1007/978-3-319-05107-9_16,

link experience with events in an excitable medium.[1] While conservatives resist, this is increasingly known as *languaging* or, more formally, human activity in which wordings play a part.[2] The terminological shift makes it harder to unzip language from lived experience. In seeking to counter verbal bias, a case study is used to how, as in seeing, languaging depends on sensorimotor contingencies (O'Regan and Noë, 2001). However, that is not the paper's focus. Rather, it aims to show that, while languaging extends embodiment, it relies on skilled action-perception. People set up synergies by fusing cultural patterns with individual history –action sets off both imagination and social perception.[3]

2 Moving Forwards Gingerly

The paper begins by considering how mind and language extend human embodiment. Using the work of O'Regan and Noë, it focuses on what it is like to have phenomenal experiences. These are traced to, not just linguistic perception, but also how language functions as what Hodges (2007) calls an action system.[4] Specifically, the paper focuses on how a person languages as, in the presence of another, he tackles a problem solving task. In a few minutes, the young man exploits languaging to shift the locus of agency between various cultural projections and his own embodiment. As in watching video, events are perceived as like a flow of pixels that *also* depicts a situation. Since this is experiential, only sensorimotor contingencies can ground what happens. Yet, unlike seeing, languaging is minded activity that prompts an actor to use *other* people's experience. As an action-system, languaging is neither organism-centred nor organism-bound. Unlike perception-systems associated with, for example, the use of seeing, hearing and touching, languaging is always under some degree of collective control. While lacking space for full discussion, the process begins as cultural values shape infant attitudes to people, events and the world's aspects (see, Cowley et al., 2004). Human use of sensorimotor contingency is normative and, strikingly, socially derived preferences arise even before birth.[5] Language skills arise

[1] Even in reading saccading prompts anticipatory action, construal and monitoring of what is seen (Järvilehto, et al., 2011). Reading too is languaging or 'activity in which wordings play a part'.

[2] Maturana applied the term to the sense-making of all species; in Becker's work, echoing Heidegger, it highlights the particular sense of utterance-acts. Human activity is thus continuous with that of other species while also associated with verbal patterns or wordings. The terminological shift is discussed by, among others, Linell (2009) and Cowley (2011b).

[3] A referee points out that this can be theorized in terms of representation. This is true. However, rather than address that view, I stress that languaging is made possible by the skilled perception of events and situations: it is akin to seeing videos as depicting episodes of life.

[4] For Hodges (and the current author), language is also a perception system and a caring system.

[5] Babies are born with preferences for the mother's voice, languages with a specific rhythmic feel and even prosodic patterns (see, Spurrett, & Cowley, 2010)

from using circumstances to manufacture and construe social affordances or, in other terms, from what Everett (2012) calls dark cultural and cognitive matter. It is important to consider how selves, or persons connect likely outcomes with self-experience. As further explained below, human actors use a language stance (Cowley, 2011a) as they play out roles in socially constituted organizations. Thus just as we are able to see films as more than changing patterns, we hear speech as less, and more, than verbal flow.[6]

3 Mind: Not Disembodied

Thirty years ago many took mind and language to be computational. Reacting against this view, the trend is to focus on embodiment. Indeed, it can be hard to understand why there was such a fuss. Rather than focus on how to explain functions (see Shapiro, 2010), I therefore make a case against reification. If not wary, like behaviourists and cognitivists, embodiment theorists may be trapped by success. In approaching constructs as hazy as 'mind' or 'experience', it is all too easy to identify the object of interest (behavior, mind) with a method for studying, say learning or cognition. This leads to muddle. Once, *behavior* was confused with what learning theory describes; later, *cognition* was ascribed to minds that compute. In embodied cognition, there is a risk of overplaying work on how bodies (and brains) regulate activity/ system-states. In O'Regan and Noë (2001), for example, perceptual modalities are said to exist 'only in the context of the interacting organism' (959). As is shown in the case study, this does *not* apply to language. Although languaging serves as a perceptual system that prompts people to heed world-side resources, it also has life-altering functions. As an action-system, languaging influences other people, one's own perception and, hence, the world perceived. In emphasising its transformative power, I highlight linguistic experience. Like a TVSS, language depends on how cultural history is keyed to cognitive biology. Just as vocalisation uses cultural patterns ('words'), the TVSS depends on manufactured parts. Both demand modes of description that capture how dynamics co-occur with phenomenal experience. By acknowledging the complexity of how language or a TVSS contribute to action, one discovers the importance of cognitive dynamics. It becomes possible to deflate verbal patterns by looking beyond both 1st person accounts and 3rd person language models. While concerted by living bodies, language is always based in social practices.

Linguistic tradition has long emphasised how language appears. The idea was formalized by Ferdinand de Saussure's evocation of a synchronic object. Language thus came to be seen as an abstract system or set of structures that, from a lay person's perspective, can be observed, but not grasped. This set an explanatory agenda for linguistics: taking word-forms for granted, the scientific challenge was to model how people (or their brains) come to identify and compose utterance-types

[6] One referee objects that, on some philosophical views, perception is not skilled. I do not claim that *all* perception by *all* species is skilled: it is enough that, say hearing Danish as Danish depends on skills in attuning to Danish ways of speaking.

from linguistic atoms (or their putative neural correlates). Such approaches draw on what can be transcribed, structure, and, having done so, covertly emphasise analogies with the seeable. Language is taken to centre on an organism, mind or brain that processes and produces speech (that is allegedly akin to its inscriptional counterparts). As in philosophy, language is separate from so-called 'language users', experience and action-perception. Opposing such views, Chomsky (1965) sought to naturalise language by appeal to mind/brain systems. It was ascribed to, not action, perception or communication, but, rather, a modular faculty that parses and constitutes strings: in later versions of his theory, the brain houses an I- language system. Language is anomalous –it uses a biological mutation or, perhaps, a spandrel.

To build an alternative based in sensorimotor contingency one has to specify language as something other than sentence-like output. The point is pressing because even challengers to classic cognitive science treat language as essentially verbal. For example, language has been traced to encodings extended by metaphor (Lakoff & Johnson, 1980), a brain that installs a serial virtual machine (Dennett, 1991) and material symbols (Clark, 2008). Often, motor experience is taken to encode verbal units in action-relevant areas of the cortex: agents are said to 'use' language because they (or brains) possess a language-system. Neural resources produce and parse –like a Latin teacher. By contrast, a view based on sensorimotor contingency rejects Chomsky's question: what do we know when we know a language? Instead, it seeks to address how languaging extends the scope of human agency. Below, this is ascribed to imagination and social perception. These link vocalization, affect, tone and bodily movement to influence thinking. By acknowledging its partly public nature, one abandons the language myth (Harris, 1981): people neither understand because they 'know' a language and nor do they 'use' utterance acts as a conduit between minds and bodies. Rather, experience of language contrasts with our many ways of using of proposition-like entities (texts). Visible wordings are merely emblems –stylised prompts. People language as they link social action to perceiving through vocal and visible gesture. Yet, as shown below, no enactive view captures the multi-scalar complexity of language. While based in the sensorimotor, agents also draw on history to treat language as picture-like.

4 From Language to Languaging

Emphasis on the cognitive dynamics of language is emerging everywhere. In ways brought home clearly by Paul Thibault (2011), there is no neat divide between the linguistic and the nonlinguistic. Just as the verbal fails to reduce to the sensorimotor, the sensorimotor is quite insufficient to explain the verbal. Languaging is something that we do: like film-making it depends on complex activity that is designed to favour perception. In films, as in languaging, the traces are multimodal, evanescent and designed for human perception. Lacking space to review this distributed perspective (see, Cowley, 2011b), I focus on the danger of separating linguistic units from action, perception and experience. For ease of exposition, I do so in Maturanian terms. While only a beginning, this biologically based approach has the merit of offering a robust challenge to word-based views of language so blithely adopted by philosophers, linguists and computationalists.

While language is seen as essentially verbal, people are separated from what they say, do and understand: often, emphasis shifts to an individual, mind or brain. Oddly, linguistic autonomy is often taken for granted. Yet, in other species, the coordination of wolves, birds, fish and bacteria is *not* due to organism centred coding (Rendall et al., 2009). Like animal communication, linguistics can be rethought (see, Love, 2004; Kravchenko, 2007). With remarkable prescience, Maturana (1978) replaced code-views by appeal to structural coupling. Precursors to language dominate the natural world: In tracing language to sensorimotor contingency, prominence falls on a history of coordinating. In humans, languaging is traced to caregiver-infant interaction that, while normative, gradually self-organises around gestural patterns (both vocal and visible). Language thus centres on communities of practice. Skillful activity, including perception, is augmented as, in Maturana and Varela's (1998) terms, people orient to the orienting of others. People draw on abstract constraints: before mastering wordings, babies learn when to fall silent, how to use a spoon and what to make of social referencing. Thus, as languaging takes on a verbal aspect, it is already anchored in a baby's experience of human ways of life.[7] A history of couplings-in-a-community prompt individuals to develop dialogical minds and brains (Linell, 2007). Given repetition with variation, agents develop what Maturana calls a consensual domain. Accordingly, other people influence how they behave and, as a result, how they individuate as persons. Though a community make and hear utterances in similar ways (using the 'same' phonetic gestures), meanings are always connotational (see, Kravchenko, 2007). The interplay of languaging is directed –not just by manifest intentions –but also by the alternation of more and less deliberate ways of pursuing outcomes. The multi-scalar nature of language enables people to engage with each other while using community-based traditions. Since human dynamics are normative they ground complex, emotion-ridden sense making; everything that we do with languaging is highly skilled.

If sensorimotor contingencies offer anything to the language sciences, it will challenge the view that writing/speaking is a matter of intentionally putting meanings into words or, indeed, using material vehicles to grasp verbal content. In Maturana's (1978) terms, speaking and writing belong to different cognitive domains. However, it is important not to overstate: experience with these domains contributes not only to new kinds of skill but also to the development of artifacts, institutions and community based modes of life.

5 A Case of Problem Solving

In tracing languaging to sensorimotor contingencies, I focus on a case where it serves to explore the world. My exposition aims to shows that, while the young man's

[7] The enactivist faces the 'problem of organizational closure'. On this view, the world is only be a source of perturbations –nothing new can enter into the world of a coupled system.

thinking is based in sensorimotor activity, his actions reach into a social domain that lies beyond body-world interaction. Just as in, say, watching a film, the air cadet draws on what he remembers, feels and imagines. At certain moments, he depends on speaking aloud while using a model to address a well-defined problem. He links what is, at any moment, being perceived with what is –and has just been –said and done. In offering a description, I return to a case study (see Cowley and Nash, 2013) of how 'Billy' solves the river problem

The problem has been studied under names such a missionaries and cannibals as well as hobbits and orcs (e.g. Jeffries et al., 1977; Knowles and Delaney, 2005). Participants are responsible for 6 parties who want to use a raft to cross a river. In approaching the task, only two parties can be placed on the raft at once. Further, in each crossing, one of the parties must row. This sets up a logical puzzle arises in that, if 'bad' guys (Ps) outnumber 'good' ones (As), they attack. For participant and psychologist alike, this constitutes failure. In the study sketched below, air-cadets chose to label the 'bad' guys pongos. The version allows an 11 move solution (1) send over A&P; (2) leave P (bring back A); (3) send over P&P; (4) bring back P; (5) send over A&A, (6) bring back A&P; (7) send A&A; (8) bring back P (9) send over P&P, (10) Bring back P; (11) Send over P&P. Any other move (or pair of moves) is banned. Participants worked under an officer's watchful eye and, in the relevant condition, used a paper maché model, raft and toy-soldiers (see Figure 1). They were told that they would receive no help and, beyond that, given no instructions about how to proceed. Next, therefore, I focus on Billy's performance. The particular case was chosen for two main reasons: (a) the cadet was one of few solvers; (b) he chose to speak about what he was doing.

Fig. 1. Billy with the model

OFFICER: ”Good”	
1. Erm	
2. so we need to send the guy (P1)	1. send over P&A;
3. my guy comes back	2. bring A back;
4. And I get another pongo erm here	
5. No	3. send over P&P;
6. So this guy over there (P2)	
7. Leave him there erh	4. bring back P;
8. And then if I take #	
9. This guy (P3) over there	Tries impossible move 1
10. I'll be attacked	
11. Should I divide my forces	
12. Or keep them together	
13. Erm # eh	
14. Second trip I'll be adding reinforcements	
15. Uhm so	
16. Keep this guy here, move the pongo (P3) over, and	Repeats impossible move 1
OFFICER: "Can't do that, cos when you get to the other bank your'e outnumbered"	
17. And uh on there	Sees impossible move 2
18. He's outnumbered	
19. Start again	Starts again
20. OK one pongo (p1)	1. send over P&A ;
21. Guy come back	2. bring A back;
22. Equal	
23. I count them equal	Impossible move 3 ??
OFFICER: Yeah you have them equal but it's not equal is it? (i.e. if you send A C)	
24. Hmmm # Ok # So	
25. If I get	
26. The pongo to drive the raft	
27. And this guy over there (P2)	3. send over P&P;
28. Equal	
29. And come back	4. bring back P;
30. And then this guy (AC1) over here	impossible move 1
31. Pongo come back	
32. This guy's outnumbered Let's back…back a step	

33. I have one pongo there	Back to 4 (leave P)
34. Which is equal	
35. Back a step	
36. It's hard eh	
OFFICER: "Two minutes left"	
37. So if I take a cadet over	
38. This guy's outnumbered	
39. If I take a pongo over	
40. my guy there's outnumbered	
######	
41. If I take my two air cadets across	
42. Then he's outnumbered there	
######	
43. Cadet goes over	
44. And he comes back, he's outnumbered	
45. Pongo come back	
46. So if I take my two air cadets over	5. send over AA
47. And I change them	6. send back AP;
48. this guy for a pongo	
49. Who gets back	
50. Take this guy over	7. send over AA;
51. Drop him off	8. Bring back P;
52. Take him back	9. Send over PP.
53. And then	10. Bring back P
54. Over there	11. Send over PP

During the first three minutes (lines 1-37) Billy is primarily concerned with practical understanding of the constraints. Then, hearing that time is getting short, he verbalizes his grasp of the situation (38-41). Having done so, he repeats himself (42-45) and, in 46, has a quasi-insight. The task is then completed in the time it takes to say 47-54 as he enacts what was meant by *Pongo come back*.

The reader is now advised to consult the transcript (see Box A). Billy continuously talks as he addresses the logical structure of the puzzle that can be traced to 11 moves. However, as the record of speech shows, he does not experience its logic. Initially, the puzzle seems straightforward: he gets two pongos to the far bank, leaving 3 cadets and one pongo on the near one. After move 4, things get hard. To capture this experience, it is called reality checkpoint (RC). Unlike many, Billy does not give up at this point. But, when he sees moves 5-6, he experiences what feels like an insight; the problem becomes easy. Billy's progress is shown on the time-chart of Figure 2 below. Having seen move 6 after 230 seconds, he then solves in only 20 seconds (making 7 moves). As explained below, special weight is given to periods marked by black bars. However, the reader should also note both errors made (bars crossing the X axis) and moves associated with numbered arrows (below the time-line). Where labeled ‘co’, the officer makes a correction; in ‘see’ cases Billy notices his own mistake (A ‘planned’ error is shown with a dotted line). It is striking that, after about 120 seconds, Billy chooses to restart. The last point to note is that he completes after attaining a quasi-insight –he ‘sees’ a solution after more than three minutes. Having seen what is the solution (without knowing that it is), he completes the task in a matter of a few seconds.

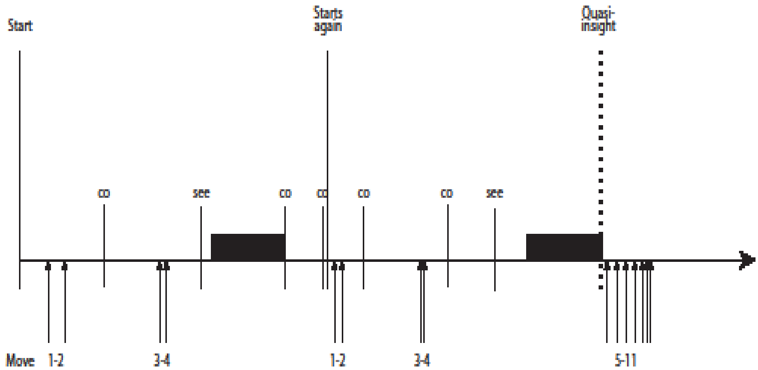

Fig. 2. Time chart showing Billy’s progress across the11 moves

6 Decisions, Decisions

Problem solving research typically focuses on how subjects find their way through an abstract problem space. However, placed in a concrete and imaginable setting, one can scrutinise what happens (see, Cowley & Nash, 2013). In Billy’s case, having been presented with the puzzle, he takes 210 seconds to grasp the problem. In so doing he comes twice to reality checkpoint (RC) and, having recognized the challenge (by making the same error twice and then testing it), he is spurred to think hypothetically. On the first occasion, he follows procedure (utterances 11-16 as shown in the first black bar). When he later returns to RC, he again addresses it in a procedural way (as shown on a black bar before the quasi-insight). However, this time, he frames the setting by opening up action-perception accompanied by utterances 43- 44 (### is a long silence):

43. Cadet goes over
44. Comes back he’s out numbered
 #################

As Billy has has just said the same thing (viz. So if I take a cadet over/ This guy’s outnumbered), these utterances are repetitive. What he has just said anticipates movements and active perception. In other words, having come up with a frame (i.e. by vocalising that he cannot send a cadet over), he tests out what he has just heard (he cannot take a cadet over).

> "Instead of being overwhelmed (or starting over), Billy faces up to reality-checkpoint. He puckers his lips and places his tongue in front of his lower teeth in a display (and, perhaps, a reminder) of motivation. Billy looks at the bank on his left and moves his hand away from where the raft is. As nothing new is visible, he seeks affordance potential in the 2 visible pongos. Next, his gaze returns to the raft as he falls motionless. Having remained still, during the silence, he then looks back across the river, stops, and then back at the raft. As he does so, he readies his hand to take up a toy soldier: however, he does not. Rather, saying, "Cadet goes over" his hand serves a proxy for acting in synchrony with, "comes back, he's outnumbered". As he does so, he shows a hint of a smile– as if having saved the cadet from the pongos. While a repetition of [41-42] (as suggested by the abbreviated form of "cadet goes over"), he checks that the situation is as verbalised. This repetition is […] contextualized by actively perceiving the world. Without being explicit, he grasps that has to do something different with the two cadets … (i.e. take one man over and bring a pongo back). Knowing that there is something more gives rise to silence and stillness (for 1600 ms) as he gazes at the two pongos on the far bank… (Cowley & Nash, 2013).

What Billy *does* constitutes the thinking. Conscious experience arises in using the model while imagining what seems impossible. The language is dominated by the said –phonetic gestures –and, thus, a fact (viz. if I take a cadet over/ This guy's outnumbered). While largely constituted by sensorimotor activity, Billy's thinking is irreducible to contingent movement. Rather, he relies on what he had said to imagine the situation and, by enacting it, explores its apparent impossibility. As so often in problem finding, he relies on stress in complex sense-making that culminates in motionless silence. Thus, after (44), Billy spends 1600 milliseconds staring at the pongos. Suddenly, a a striking gesture is integrated with saying:

45. Pongo come back

Although showing no sign of recognising this as the solution, it *sounds* like an insight (for detail, see Cowley and Nash, 2013). Billy uses a rapid hand-movement that is synchronized with 'pongo come back' (with a high pitched head).[8] While the tone is notable what is more striking is the prosodic echo of 'cadet goes over' (also spoken with a high head). He has the correct way of going on. Though he has taken

[8] In other words, the first syllable of 'pongo' is given some prominence by a rise in pitch; it is, however, not the nucleus of the syllable. Phonetically, this is highly marked.

230 seconds, he can now solve in 20. The utterance-act anticipates what he is about to do. Though compressed, it amounts to saying: "If I take a cadet over and swap him with a pongo who comes back, I might be able to solve the puzzle". However, while phonetic gesture meshes with visible movement, both seem to come from nowhere. How does staring at pongos prompt this act? The question can only be understood by tracking back to what Billy had previously done.

The steps which prompt his to imagine the solution are no mystery. In brief, almost three minutes earlier (210 secs), Billy had been stuck at reality checkpoint. Having made the error shown before the black block before restart (Figure 2), Billy faced up to problem. He 'saw' that if he took a cadet to the other bank he would be attacked. So, detaching himself from events, he echoes his own training:

11. Should I divide my forces
12. Or keep them together
13. Erm # eh
14. Second trip I'll be adding reinforcements
15. Uhm so
16. Keep this guy here, move the pongo over, and

    OFFICER: "Can't do that, cos when you get to the other bank your'e outnumbered"

Repeated blindly, his training leads him to 'forget' his practical lesson. In spite of having just seen that he could not take over another pongo (or 'divide his forces'), his utterance induces him to repeat his error. No doubt, this compounds the stress. This too shapes the solution-probing that arises when Billy returns to reality checkpoint (RT). For a cognitive psychologist, the situation is identical (he is in the same abstract problem space). However, this is not Billy's experience. What he does is fuelled by feeling as, once again, he distances himself from the task by rehearsing his experience so far.

37. So if I take a cadet over
38. This guy's outnumbered
39. If I take a pongo over
40. my guy there's outnumbered ######
41. If I take my two air cadets across
42. Then he's outnumbered there ######

Using working memory (or similar) Billy keeps the facts in mind. As he now knows (as long as he thinks one step ahead), he cannot take a single cadet over. Further, he cannot take a pongo over and nor can two cadets go across the river. Languaging makes Billy's experience picture-like: holding 'facts' in mind, he can imagine or hypothesize 'impossible' solutions. These, of course, *cannot* be reduced to the sensorimotor. Thus, while overlooking the rider (viz. what he says applies if, and only if, he chooses not to think 2 moves ahead), he pinpoints the issues by narrowing the problem space to a single option. Just as with syllogistic logic, he reframes.

Without being explicit, *what is said* becomes a proxy for the world.[9] As shown by the quasi-insight, he prompts himself to see that, using the move after next (pongo come back), the problem can be solved. He *can* take two cadets across –if and only if, in bringing the raft back, he swaps one with a pongo. This is the core point: the duality of languaging (and language) depends on using sensorimotor activity as a normative resource that contributes to an individual history of perceiving acts as wordings.

7 Languaging and Sensorimotor Contingency

The case study's first lesson is that neither problem-solving, action, nor language rely exclusively on intracranial events. Much depends on what Billy does and says: in other words, there is reason *not* to conceptualise language as based on putting (inner) meanings into words. Even in this situation, like structural coupling, utterance-acts alter perception. Specifically, they prompt Billy to reframe experience by articulating what he believes cannot be done. This is other-oriented: first, to a remarkable extent, both the officer and an analyst can track how he feels, thinks and acts. Second, Billy himself uses what he says to reformat his experience. Third, and most crucially, Billy uses languaging as a mode of action. By performing as a well-trained air cadet, he exploits the non-present. He draws effectively on the picture-function of language: while, at times, he is fully absorbed (especially during 43-45) this does not apply during periods shown in black (in Figure 2). When at reality checkpoint, he twice takes his distance by using (among other things), training, logic, distaste for pongos and a desire to succeed. He links the event flow with *imagining*. He shifts his agency (and attention) between involved and detached engagement with the situation.

While a philosopher might offer an explanation, my aim is descriptive. For my purposes it is enough that, just as we see pictures as pictures of something, Billy hears languaging as languaging about something. As a trained and educated air cadet, he uses wordings to unearth complexity. Though based in social contingency, movement of the articulators gives rise to imaginative experience. This connects speech to 'rule-based' knowledge or, in other terms, allows self-display to be used in construing what can be perceived. Although wordings are *abstracta* with a cultural history, they need not be 'realised'. While an attempt at explanation might invoke 'mental representations', this would be like arguing that the dark side of the moon is made of blue cheese. Indeed, most of Billy's thinking–including generating a quasi-insight – *is* sensorimotor activity. No representational model can explain, for example, how HIS tongue movements can sustain attention. However, Billy also uses languaging in its picture-function. He takes a language stance (see, Cowley, 2011a) by regarding his

[9] Language enables us to do something like looking at a picture: it prompts us to perceive arrangements between entities or, in Gibson's (1979) terms, to attend to not only the invariants *of* the picture but also invariants *in* the picture (see Cowley, 2011a). While this can be theorized in terms of representations, there is no need to do so. Indeed, Wittgenstein's reaction to the Tractatus can be read as seeking to clarify this 'picture-function' (one that arises, roughly, when a person chooses to perceive languaging, or its traces, as about relations between entities/classes of entity).

utterance acts *as* ‘about the situation (i.e. as words with meanings). He depends on the talk of those who came before and is, in this sense, his act is irreducible to the sensorimotor (or phonetic gestures). Of course, to show that there are bursts of speech that involve hypothetical thinking –and the use of procedures –does not show that language lacks a sensorimotor grounding. What it brings to the fore is that so-called ‘use’ or ‘knowledge’ of language is intrinsic to bodily activity: languaging is no more explained by representations than bursts of utterance-activity reduce to the sensorimotor. Since language serves action and perception (not to mention other functions), Billy need not attend to how exactly he moves and articulates as he probes the model. Rather, sensorimotor experience informs how he lives the moment. In Everett’s (2012) terms, he draws on dark cognition as circumstances suffuse the utterances with what he experiences as the particular sense of the events.

Close up, there are major contrasts between languaging and seeing. Whereas utterances 43-44 are perception-like in giving actions transparency, this observation lacks general application. Language is no modality that ‘exists only in the context of an interacting organism’ (O’Regan and Noë, 2001: 959). Rather, while having a sensorimotor basis, it also allows people to use languaging as an action system that re-evokes cultural resources. It depends on virtual patterns or future attractors that influence the play of phonetic gestures. This, indeed, shapes Billy’s quasi-insight. Quite literally, he uses 1600 ms of intensive looking to go beyond the information given. Drawing on frustrations, he sees that he can take two cadets across –provided that he swaps one with a pongo. However, this is *not* what he says (he says ‘pongo come back’): he is not bound by the denotations of word-forms. Rather, striking formatting marks an (unspoken) ‘idea’. While less dramatic, the other cases also show that languaging can alter a perceived (social) situation. In 11-16 and 37-42, Billy actively distances himself from events or, in the haunting phrase, uses dark cultural matter. He draws on procedures –one from air-force logic and one based on skills in reasoning. Skilled linguistic action gives him distance from what he sees. Other peoples’ language prompt the cadet’s affordance-making. In other words, Billy also acts to create a lived situation. Although seeing affects external memory, skilled languaging reshapes a perceived world. This arises since, as Bakhtin (1984) emphasised, languaging fuses peoples’ experience. Thus while verbal patterns can be said to evoke a consensual domain, the case study shows that the concept stands in need of very substantial development.

8 Languaging, Contingency and Verbal Patterns

In objecting that linguistics has been transfixed by verbal bias, I used Maturana’s work as an antidote. As far as it goes, it is good enough: not only can languaging be traced to biology but what he deems ‘structural coupling’ engenders perceptual experience. In this sense, language is rightly compared with seeing. As applied to the case study, the view helps clarify how language contributes to testing the ‘truth’ of utterances 43-44. However, as an action system, languaging is *more than* structural coupling. Billy uses learned procedures to bring forth a new perspective. He conducts himself as an air-cadet who seeks to solve a psychological problem. Not only is this irreducible to a history of structural couplings but, just as clearly, it is irreducible to an organism-constructed consensual domain.[10] Although akin to seeing in that language alters our grasp of circumstances, language also has the power to transform what the world affords. In saying 'pongo come back' (or 11-16; 37-42), Billy's speaking prompts him to discover a viable solution. While based in individual history, the act of utterance is conceptually constrained. Both languaging and language combine properties of a sensory modality with the heuristic use of conceptual resources.

Linguistic coupling functions in different time-scales. In drawing the analogy with seeing, in viewing language as a perceptual system, the focus falls on the specious present. If we attend to acts of utterance, language has transparency. However, Billy is not wholly dependent on real-time dynamics. In seeking a solution, he skillfully makes reciprocal connections between experience and the said. Just as in O'Regan and Noë's (2001) account, in language too, there are degrees of awareness. At times, people speak in idiosyncratic ways (e.g. pongo come back) and, at others, they rely on conventions and modes of life. Used as an action-system, what is said takes on a striking prominence: thus, in 11-16, Billy uses it to re-evoke his training (with some accuracy) and, in 37-42, he adeptly summarises his understanding of the task. Most strikingly, in 45, he correctly anticipates the solution. In all such cases, he treats his utterances as utterances of something or, in other terms, he uses skills with taking a language stance. At such times, like a TVSS, language is at once both a cultural product and an action-perception system. While drawing on sensorimotor contingencies, people use history of exploiting language about language (and social 'reality'). Linguistic reflexivity permeates the specious present. Billy uses this when, in 37-42, he renders nonce experience explicit; using the language stance, he imperfectly grasps the puzzle's logical constraints. While lacking space to pursue the view, the language stance undergirds much languaging in children over the age of about 2 and, later, is enhanced by experience of literacy. Other perceptual systems lack any clear parallel (one cannot see about seeing).

Languaging demands a conception of language (and, indeed, languages). This is because, as people language, rich sensorimotor dynamics come to be heard as iterating verbal patterns. Language is, at once, dynamic and symbolic. Indeed, Billy uses this strange duality to shift between modes of acting. While his doings centre on organism-environment relations (and conscious experience), he also draws on community life. Using procedures, he can think like an air cadet who is on a training exercise or, indeed, like a leader who argues logically. In the puzzle, the leadership role is more effective. Billy combines the sense of utterance events with meaning potential. This enables him, if he chooses, to give them a lasting sense. Moments of languaging come to be treated as indicative (11-16), factual (37-42) or insightful (45). Although derived from a history of contingencies, these hearings echo linguistic

[10] As Brier (2008) points out, this is usually construed as actualizing meaning that is already there. In fact, people develop conceptual schemes to stabilize self-reference over time; they also use other people's experience to develop action-systems.

coupling in a collective domain. Just as is in experimental work on how conventions arise, the patterns are not semantically neutral (Mills, 2013). They depend on what, in simplified settings, Mills (ibid.) calls the tacit negotiation that shapes joint experience as people adopt increasingly complementary modes of action.

9 Beyond Computer Metaphors

Like seeing, languaging can activate knowing. It is almost certainly grounded in a history of sensorimotor contingencies. Indeed, were this not so, one could not expect language and perception to intermesh in subtle ways. Like movement, the verbal aspect of language serves in directing attention. Its grabbiness influences experience. As a result, people become sensitive to individual and collective influences on how and what they perceive. They draw on communities of practice that influence how they see objects, live situations and experience various kinds of events. We develop, or act as if we possess, perceptual imagination. Perhaps all of this would be obvious – were it not for the verbal focus of linguistic and philosophical tradition. Countering, the case study traces linguistic experience, hearing, and remembering to acting. Billy manufactures ways of thinking and perceiving as he goes beyond the information given. Languaging while looking leads to a quasi-insight: he makes explicit what had been hidden –saying ‘pongo come back’ is redolent with meaning.

Perceptual modalities have evolved often and independently in many species. While Maturana is surely right that all languaging has a common history, it also appears to be *sui generis*. Wordings give people actional powers: as the red queen remarked, one can have several impossible thoughts before breakfast. For this reason, language depends on perception: thinking must be constrained by ‘reality’ (in social and material aspects). Perceptual and actional experience thus drives what Gibson (1979) calls the education of attention. Although we hear every utterance uniquely, action occurs under collective constraints. Not only does this normative dimension emerge from the latter stage of embryonic development but it ensures that human life develops in ways that demand accountability. For this reason, then, linguistic activity is richer than sensorimotor experience (or structural coupling). It is best conceptualized as sense-saturated coordination or *interactivity* (see, Steffensen, 2013; Cowley & Vallée-Tourangeau, 2013). Not only does this link perception, action and experience but it is imbues experience with normative concerns. Much depends on learning to say things and consider things. The taste of wine or one’s reaction to quality writings is mediated by sensorimotor knowledge, personal experience and a community’s conceptual products. The duality of language give a hearer experience of linking the sensorimotor to community patterns. Wordings carry hidden information or, if one prefers, dark cognitive and cultural matter. Sensing their power, linguists wrongly identified languages with sets of utterance-types (see, Bloomfield, 1933; Harris, 1951). By leaving out sensorimotor dynamics, Chomsky’s (1965) return to mentalism was inevitable. Billy’s case study thus offers two simple warnings. First, just as language does not reduce to words, it does not reduce to sensorimotor contingency. Languaging is metabolically-based use of pattern that constrains experience: it grounds skills in engaging with the world. Arguably, the second lesson is even more fundamental. Billy shows that an object of enquiry (whatever we take that to be) must be separate from an investigator's favoured methods. Language extends the sensorimotor: people use training to lock onto community-based procedures. Often these presuppose a language stance or skills based on hearing utterances as utterances of something. In Sartre's (1945) hell, human understanding is largely derived from *other* people.

References

Bakhtin, M.M.: Problems of Dostoyevsky's poetics. University of Minneapolis Press, Minneapolis (1984), Emerson, C. (ed.)

Bloomfield, L.: Language. Holt, Rinehart and Winston, New York (1933)

Brier, S.: Cybersemiotics: why information is not enough! University of Toronto Press, Toronto (2008)

Browman, C., Goldstein, L.: Toward an articulatory phonology. Phonology Yearbook 3(1), 219–252 (1986)

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

Chomsky, N.: Aspects of the Theory of Syntax. MIT Press, Cambridge (1965)

Cowley, S.J.: Taking a language stance. Ecological Psychology 23(3), 185–209 (2011a)

Cowley, S.J.: Distributed language. John Benjamins, Amsterdam (2011b)

Cowley, S.J., Moodley, S., Fiori-Cowley, A.: Grounding signs of culture: primary intersubjectivity in social semiosis. Mind, Culture and Activity 11(2), 109–132 (2004)

Cowley, S.J., Vallée-Tourangeau, F.: Systemic cognition: Human artifice in life and language. In: Cowley, S.J., Vallée-Tourangeau, F. (eds.) Cognition Beyond the Brain: Interactivity, Computation and Human Artifice. Springer, Berlin (2013)

Cowley, S.J., Nash, L.: Language, interactivity and solution probing: repetition without repetition. Adaptive Behavior 21(3), 187–198 (2013)

Daniel, D.: Consciousness explained. Little Brown, Boston (1991)

Everett, D.L.: Language: the cultural tool. Profile Books, London (2012)

Fusaroli, R., Tylén, K.: Carving language for social coordination: A dynamical approach. Interaction Studies 13(1), 103–124 (2012)

Gibson, J.J.: The senses as perceptual systems. Houghton Mifflin, Boston (1966)

Gibson, J.J.: The ecological approach to visual perception. Houghton Mifflin, Boston (1979)

Harris, Z.: Methods in structural linguistics. University of Chicago Press, Chicago (1951)

Harris, R.: The language myth. Duckworth, Oxford (1981)

Heidegger, M.: On the way to language. Harper & Row, New York (1971); Originally published, Unterwegs zur Sprache (1959)

Hodges, B.H.: Good prospects: Ecological and social perspectives on conforming, creating, and caring in conversation. Language Sciences 29(5), 584–604 (2007)

Järvilehto, T., Nurkkala, V.M., Koskela, K.: The role of anticipation in reading. Pragmatics & Cognition 17(3), 509–526 (2009)

Jeffries, R., Polson, P.G., Razran, L., Atwood, M.E.: A process model for missionaries-cannibals and other river-crossing problems. Cognitive Psychology 9(4), 412–440 (1977)

Knowles, M.E., Delaney, P.F.: ‘Lasting reductions in illegal moves following an increase in their cost: evidence from river-crossing problems. Journal of Experimental Psychology: Learning, Memory, and Cognition 31(4), 670 (2005)

Kravchenko, A.V.: Essential properties of language, or, why language is not a code. Language Sciences 29(5), 650–671 (2007)

Lakoff, G.J., Johnson, M.: Metaphors we live by. University of Chicago Press, Chicago (1980)

Linell, P.: The written language bias in linguistics: Its nature, origins and transformations, vol. 5. Psychology Press, Brighton (2005)

Linell, P.: Dialogicality in languages, minds and brains: is there a convergence between dialogism and neuro-biology? Language Sciences 29(5), 605–620 (2007)

Linell, P.: Rethinking language, mind and world dialogically: Interactional and contextual theories of sense making. Information Age Publishing, Charlotte (2009)

Love, N.: Cognition and the language myth. Language Sciences 26(6), 525–544 (2004)

Maturana, H.: Biology of language: the epistemology of reality. In: Miller, G.A., Lenneberg, E. (eds.) Psychology and Biology of Language and Thought: Essays in Honour of E-Lenneberg, pp. 27–63. Academic Press, New York (1978)

McNeill, D.: Hand and mind: what gestures reveal about thought. University of Chicago Press, Chicago (1992)

Merleau-Ponty, M.: Phenomenology of perception. Routledge, London (1996)

Mills, G.J.: Dialogue in joint activity: complementarity, convergence and conventionalization. New Ideas in Psychology. Corrected proof (April 24, 2013)

O’Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24(5), 939–972 (2001)

Port, R.F., Leary, A.P.: Against formal phonology. Language 81(4), 927–964 (2005)

Rączaszek-Leonardi, J., Kelso, J.A.S.: Reconciling symbolic and dynamic aspects of language: Toward a dynamic psycholinguistics. New Ideas in Psychology 26(2), 193–207 (2008)

Rendall, D., Owren, M.J., Ryan, M.J.: What do animal signals mean? Animal Behaviour 78(2), 233–240 (2009)

Sartre, J.P.: Huis clos. Gallimard, Paris (1945)

Shapiro, L.: Embodied cognition. Routledge, London (2010)

Steffensen, S.V.: Human interactivity: problem-solving, solution-probing and verbal patterns in the wild. In: Cowley, S.J., Vallée-Tourangeau, F. (eds.) Cognition Beyond the Brain: Interactivity, Computation and Human Artifice, pp. 195–221. Springer, Berlin (2013)

Spurrett, D., Cowley, S.J.: The extended infant: utterance-activity and distributed cognition. In: Menary, R. (ed.) The Extended Mind, pp. 275–323. MIT Press, Cambridge (2010)

Thibault, P.J.: First-Order Languaging Dynamics and Second-Order Language: The Distributed Language View. Ecological Psychology 23(3), 210–245 (2011)

Wittgenstein, L.: Philosophical investigations, 2nd edn. Blackwell, Oxford (1965)

Wittgenstein, L.: On certainty. Blackwell, Oxford (1980)

Author Index