

Phenomenological constraints: a problem for radical enactivism

Roberts, Michael

DOI:

[10.1007/s11097-017-9511-5](https://doi.org/10.1007/s11097-017-9511-5)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Roberts, M 2017, 'Phenomenological constraints: a problem for radical enactivism', *Phenomenology and the Cognitive Sciences*. <https://doi.org/10.1007/s11097-017-9511-5>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Phenomenological constraints: a problem for radical enactivism

Michael Roberts¹

© The Author(s) 2017. This article is an open access publication

Abstract This paper does two things. Firstly, it clarifies the way that phenomenological data is meant to constrain cognitive science according to enactivist thinkers. Secondly, it points to inconsistencies in the ‘Radical Enactivist’ handling of this issue, so as to explicate the commitments that enactivists need to make in order to tackle the explanatory gap. I begin by sketching the basic features of enactivism in sections 1–2, focusing upon enactive accounts of perception. I suggest that enactivist ideas here rely heavily upon the endorsement of a particular explanatory constraint that I call the *structural resemblance constraint* (SRC), according to which the structure of our phenomenology ought to be mirrored in our cognitive science. Sections 3–5 delineate the nature of, and commitment to, SRC amongst enactivists, showing SRC’s warrant and implications. The paper then turns to Hutto and Myin’s (2013) handling of SRC in sections 6–7, highlighting irregularities within their programme for Radical Enactivism on this issue. Despite seeming to favour SRC, I argue that Radical Enactivism’s purported compatibility with the narrow (brain-bound) supervenience of perceptual experience is in fact inconsistent with SRC, given Hutto and Myin’s phenomenological commitments. I argue that enactivists more broadly ought to resist such a concessionary position if they wish to tackle the explanatory gap, for it is primarily the abidance to SRC that ensures progress is made here. Section 8 then concludes the paper with a series of open questions to enactivists, inviting further justification of the manner in which they apply SRC.

Keywords Perception · Enactivism · Explanation · Phenomenology · Explanatory gap

✉ Michael Roberts
MDR465@bham.ac.uk

¹ Department of Philosophy, Theology and Religion, University of Birmingham, B15 2TT, Birmingham, UK

1 Introductory comments

The last two decades have seen steadily growing interest in the enactive approach within cognitive science and philosophy (e.g., Colombetti 2014; Noë 2004; Stewart 2010/2014; Thompson 2007). Originally articulated as an approach to understanding and studying the mind, early enactivism stressed that cognition is best conceived as a form of *embodied action* (Varela et al. 1991/1993, p. 172) and made two central claims: firstly, cognition is ontologically continuous with the basic life-regulation processes of living organisms (pp. 87–9); secondly, cognition’s nature is strongly dependent upon the embodiment of the organism in question, and thus the kinds of sensorimotor interaction with the environment that it can maintain (p. 173). As recently noted (De Jesus 2016), contemporary enactivist theorising has centred largely upon how best to systematise and develop the approach, leading to the crystallisation of a number of competing versions of enactivism that afford varying degrees of significance to the claims set out in *The Embodied Mind* (Varela et al. 1991/1993).¹ This diversity has understandably been accompanied by controversy over the authenticity of these different approaches to early enactivist ideas (see Di Paolo et al. 2010/2014).

In this paper, however, I am not much concerned with these theoretical details. I am interested rather in the *means* by which the basic enactivist picture was originally derived and the broad repercussions that this has for the future development of enactivist thought. Specifically, I shall be concerned with the suggestion that enactivism is a thesis arrived at principally through *phenomenological reflection* (Thompson and Cosmelli 2011, p. 165).

The founders of enactivism note that their ideas were chiefly derived from the phenomenological considerations and analyses of Merleau-Ponty (1945/1962), which stress the embodied and interactive qualities of experience, and they mark *The Embodied Mind* (Varela et al. 1991/1993) as an extension of his research project (p. xv). This text thus set out a bare-bones model for how the cognitive sciences could proceed in a way that was informed by Merleau-Ponty’s phenomenological insights, namely, one which focused cognitive science’s empirical lens upon the appropriate horizons: the embodied interactions maintained between organisms and their environments. Applying this line of thinking to *perceptual* cognition, Thompson and Cosmelli note that they ‘use phenomenological considerations about perceptual experience to constrain how [they] think about the subpersonal mechanisms of perception’ (2011, p. 165) and they recommend that cognitive science more broadly take this approach, interrogating subpersonal mechanisms extending across brain, body and environment.²

This paper will be concerned firstly with two important questions that arise in response to this strategy, especially amongst those unfamiliar with the continental literature from which enactivism springs: (1) Why ought we to think that phenomenological considerations should be able to influence our cognitive science? (2) In precisely what manner should this occur? An initial aim of this paper is to answer these in detail,

¹ The main contenders here are *autopoietic enactivism* (e.g., Thompson 2007; Di Paolo et al. 2010/2014; Colombetti 2014), *sensorimotor enactivism* (e.g., O’Regan and Noë 2001; O’Regan 2011; Degenaar and O’Regan 2015; c.f. Hurley 1998) and *radical enactivism* (Hutto and Myin 2013). These terms are due to Hutto and Myin (2013). See Hutto and Myin (2013, pp. 23–36) for a good outline of the principal differences between these contenders.

² Those researchers who have taken up this model now make up what is called *enactive cognitive science*. See Di Paolo et al. (2010/2014) for more on this.

for I believe that the motivation for Thompson and Cosmelli's stated explanatory strategy remains unfortunately opaque in the literature, particularly to thinkers working exclusively in the analytic tradition, who have inherited a long-standing suspicion over the utility of first-person considerations when it comes to theorising about the mind.³ Without clarity here, enactivists have been left open to sustained critique over the seeming powerful influence that it affords to phenomenological data (e.g. Rupert 2009, 2015) and a persistent misunderstanding of basic enactivist motivations.

I will attempt to rectify this situation in the first half of this paper, unpacking the reasoning behind Thompson and Cosmelli's approach at length. I will begin by sketching some minimal details about enactivism in section 2, before moving on in section 3 to introduce a particular explanatory constraint that the majority of enactivists have picked up, which I shall argue can help us to understand why enactivists make the claims that they do. I call this the *structural resemblance constraint* (SRC). I shall outline the justification for SRC in section 4, where I will mark SRC as an attempt to motivate identity claims between the mental and the physical, and as itself supported by standalone considerations from the philosophy of science which are familiar to analytic thinkers. In section 5, I shall then crystallise the manner in which enactivists have made use of SRC in order to arrive at their views on perception.

The second aim of this paper will then be critical, for the reliance upon SRC within the enactivist camp is not without its problems. In sections 6–7, I will reveal this difficulty in the “Radical Enactivism” of Hutto and Myin (2013). Despite implicitly supporting SRC as a means of motivating identity claims, I believe that their programme for Radical Enactivism contains claims which threaten to undermine SRC completely. The problem here is Radical Enactivism's stated consistency with the narrow (brain-bound) supervenience of perceptual experience. I shall argue that, given Hutto and Myin's phenomenological commitments, SRC is inconsistent with narrow supervenience, and that one of the two must be rejected.

Focusing on the case of Radical Enactivism this way will allow me to arrive at a broader and more significant point. I will outline that the endorsement of SRC is critical to the claim that enactivism can tackle the explanatory gap. As such, I will argue that enactivists must avoid being concessionary on the issue of supervenience, and *reject* the possibility of (perception's) narrow supervenience, if they wish to preserve enactivism's progress against the explanatory gap and promote its further bridging. The remainder of the paper (section 8) will then offer some questions for enactivists concerning their application of SRC.

2 Basic features of enactivism

To narrow the subject matter, I focus here upon enactivism's construal of a single subset of (human) cognition: perception.⁴ Speaking very broadly, enactivists claim that

³ For a good recent example of such suspicion, see Schwitzgebel (2008). For a diagnosis of how this suspicion came to be commonplace, see Hatfield (2005).

⁴ My treatment of perception takes it to be an essentially *conscious* process. Though the term is also used to refer to unconscious aspects of sense-making, it here concerns only those mental states/processes that there is “something it is like” for the subject to undergo. In the remainder of the paper, I will thus equivocate between perception and perceptual experience, as is common amongst enactivists. I will also use the term ‘process’ rather than ‘state’, given the enactivist understanding of perception as something which takes time to happen.

perception is a way of acting in the world, wherein a meaningful perceptual world is constituted at least partly by the specific activities the organism performs (Noë 2004; Di Paolo et al. 2010/2014, pp. 39–40). Enactivists suggest that perception consists not merely in *receiving* the world, but also in *choosing* which parts of the world one will be affected by, which aspects of the world are brought into an organism's sensorimotor loop (Stewart 2010/2014, p. 3), shown in Fig. 1.

To explain, the embodied activity of the organism (O) is said to determine the aspects of the environment (E) that will impinge upon it, and thus its next sensory inputs, while those sensory inputs are used to guide the motor actions that are needed to sustain things within this loop, themselves required to maintain the organism's viability (Stewart 2010/2014, pp. 1–4; Thompson 2007). A central claim of enactivism is therefore that *both* of these sensory and motor processes go into perception.

For present purposes, the details of this process are insignificant. We need only observe that the circular phenomenon outlined in Fig. 1 is christened as the '*enaction*' or '*bringing forth*' of a world (Varela et al. 1991/1993, p. 205), that is, the *allowing of something to present itself* or to show up for the organism. Given this account of perception, it is crucially important for the enactivist that one be in constant contact with the environment *out of which* a world is enacted. Accordingly, it has become common more recently to regard the enactive understanding of perception as something not merely active and embodied, but also *interactive*. Thompson and Cosmelli have thus qualified in a more recent article that 'to perceive is to be in an interactive relationship with the world' (2011, p. 165; c.f. Noë 2004; Pepper 2014; Ward 2012).

Now, this is an understanding that, as I have said, is purported to arise from phenomenological considerations about the nature of perceptual experience. I suggest that enactivist motivations to assert this understanding depend not merely upon a certain Merleau-Pontian phenomenological picture, but also upon the concomitant endorsement of a particular explanatory constraint that I call the *structural resemblance constraint* (SRC).




Fig. 1 A sensorimotor loop (arrows show direction of influence)

3 The structural resemblance constraint

SRC is a constraint upon explanations of *conscious/phenomenal* mental processes (those that there is *something it is like* to undergo) within cognitive science. Before outlining SRC, it will be useful to outline two broad distinctions that can help identify the type of explanation that SRC targets.

Firstly, we have a distinction between *horizontal explanations* and *vertical explanations* (Cummins 1983). *Horizontal explanations* seek to account for something's occurrence by proffering a sequence of distinct *preceding* events (c.f. Drayson 2012; Bermúdez 2006). For example, we might seek to explain the occurrence of a tornado by reference to a series of pressure fronts and air moisture levels converging over time. This form of explanation thus thematises *why or how something was brought about*. Contrastingly, we have *vertical explanations*, which offer some set of constituent parts and relations that collaboratively make up or constitute some phenomenon, thereby accounting for its features and causal powers (c.f. Drayson 2012, p. 2). For example, we might explain the properties and causal powers of a tornado by referencing the real-time internal rotation dynamics of the air it contains. This form of explanation thus thematises *what something is made of*.

Both of these forms of explanation are found within cognitive science. For example, a clinical psychologist might explain some emotion by citing a person's previous attitudes and negative experiences - a horizontal explanation. Contrastingly, a cognitive neuroscientist might seek to explain the same emotion by relating the set of neuronal processes and relations sufficient to constitute it - a vertical explanation.

With this distinction on the table, I can state that it is solely *vertical explanations* that concern SRC. Furthermore, it is *complete*, rather than *partial*, vertical explanations which SRC targets - those which seek to proffer the necessary *and sufficient* constituent parts and relations of some phenomenon. A second distinction helps to further clarify the kind of vertical explanations of interest here. This is a distinction between *personal-level* and *subpersonal-level* explanations.

A *personal-level* explanation will invoke that which we ordinarily ascribe to *whole persons* (e.g. beliefs, desires and experiences) in its explanatory narrative. It will speak in *personal* terms. This is the mainstay of folk psychology, where we might explain some individual's behaviour, say, by invoking their preceding desires. A *subpersonal-level* explanation, meanwhile, will invoke things that we do *not* ordinarily ascribe to *whole persons*. It will speak in *subpersonal* terms taken from the larger domain of inquiry from which the explanation emerges. For example, a neurological explanation of some mental process will speak in the terms of neurological theories, such as "neuronal activations" or "long-range synchrony", while a quantum-mechanical theory will use very different terms.

While I am concerned here solely with vertical explanations, it is perfectly possible for these to proceed at either the personal or subpersonal-level. For example, Introspectionist thinkers (e.g. Titchener 1899) have offered accounts of mental processes as being built out of simpler experiential *elements* or *atoms*, thereby giving vertical explanations in personal-level terms.⁵ However, these kinds of explanation (both vertical *and* personal-level) are not typically considered the task of contemporary

⁵ See Kriegel (2015) for a contemporary exploration of related ideas.

cognitive science, and I set these aside here too. Cognitive science is routinely taken to account for the personal-level in terms of the subpersonal-level, and it is this kind of *subpersonal vertical explanation* that SRC concerns. This clears the space to offer its full outline.

The structural resemblance constraint (SRC) An explanation of some conscious mental process X (which is vertical and subpersonal) should reveal a strong *structural resemblance* between (a) the combined constituent parts and relations that it invokes and (b) X as it is best characterised phenomenologically.

This needs unpacking. Firstly, a structural resemblance holds iff two relata possess some shared set of structural features, i.e. iff a mapping of one domain onto the other preserves at least some of its structural relations. If two relata share such structural features, it will be possible to abstract away from the two to reveal a structural description (a formal model) that is equally accurate in describing both (Bayne 2004, p. 360). The limit case of this is an *isomorphism*, where the mapping is both one-to-one and *fully* structure preserving.

Given that a set of shared structural features can be larger or smaller, structural resemblance occurs on a spectrum. For example, suppose we have three pencils, all of which possess the same overall macro-shape (all are equally long, pointed cylinders). The first two will be more strongly structurally resemblant to each other than they will be to the third, if these two are both traditional wooden pencils with a common simple internal make-up, while the third is a mechanical pencil with its own more complex internal organisation.

Secondly, as to precisely *how strong* a resemblance SRC demands, we can think of this as a question of degree; the stronger the resemblance revealed, the more satisfactory the explanation will be, having taken all other factors into account. I will return to this important issue later.

Finally, relata (b) in SRC refers only to the entity as invoked and characterised *intrinsically* by our best phenomenological accounts, rather than in terms of the causal role it plays within a larger system, or the relations it holds to other mental processes. Regarding what our best phenomenological accounts might be, and where they might come from, I will tackle these issues in subsequent sections.

Now, SRC is a weaker formulation of the structural *isomorphism* constraint picked up in the work of Lutz (2002), Metzinger (2000, pp. 67-9), Varela (1997) and Wheeler (2005, p. 133, pp. 225-236, 2013, p. 147), thus allowing for *degrees* of explanatory value to emerge, rather than proposing an all-or-nothing bargain.⁶ I also believe that SRC is implicitly endorsed in the work of a majority of self-avowed “enactivist” thinkers, as I shall go some way to demonstrate. I am in favour of such endorsement, for I believe that SRC is valuable as an explanatory constraint. However, I think the reasons for this are not always transparent in the literature

⁶ Though these thinkers speak in terms of ‘isomorphism’ rather than ‘resemblance’, I suspect that most of them would be willing to grant the value of this weaker formulation. Also worth noting here, while it is common to invoke dynamical-systems-theory (DST) as providing an appropriate kind of abstract model (a *mathematical* model, in this case) with which to capture a structural resemblance (e.g. Lutz 2002), SRC is itself more general and can allow for less technical models to be utilised. We can therefore think that a structural resemblance holds to the extent the *any* broadly-construed abstract structural model can be applied to both relata.

at present. As such, I want to offer a standalone justification for SRC here. This manner of selling SRC will show it to tap directly into more mainstream explanatory concerns in the philosophy of science and mind, thereby attempting to make enactivism more comprehensible to a broader swathe of thinkers and, simultaneously, opening up a larger store of existing philosophical material for enactivists to use in support of their views. In brief, I wish to implicate SRC as no more than a specific application of something more widely discussed: a constraint upon *reduction* in the sciences. Doing this will allow me to make two points in support of SRC: (A) SRC mirrors a specific criterion for ontological reduction in the sciences that serves to increase the *intelligibility* of making identity claims, and thus (B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes. (A) is the central point, while (B) shows the benefit of (A) in the realm of cognitive science.⁷ Let us take (A) first.

4 Reduction and identity

Before beginning, an account of reduction must be offered. A loose sense of the idea as given in the *Stanford Encyclopedia of Philosophy* will be sufficient here:

The term ‘reduction’ as used in philosophy expresses the idea that if an entity x reduces to an entity y then y is in a sense *prior to x* , is *more basic than x* , is such that x *fully depends upon it* or is *constituted by it*. Saying that x reduces to y typically implies that x is *nothing more than y* or *nothing over and above y* (van Riel and Van Gulick 2016, ‘Scientific Reduction’, para 1, emphasis in original)

What then are some reasonable constraints upon a reduction? I will look here at some recent comments on this issue in Hutto and Myin’s *Radicalizing Enactivism* (2013). By examining the criteria invoked here, which Hutto and Myin themselves endorse, we shall see that SRC offers merely a domain-specific application of one such criterion. Consequently, by interrogating the reasons given in favour of such a criterion, we shall have reasons in favour of SRC.⁸

Hutto and Myin (2013) first note that contemporary philosophy of science no longer holds inter-theoretic reductions to require strict a priori *deduction* of the truths of one domain from the truths of another (p. 175), as is sometimes stipulated concerning the mental and its relation to the physical (e.g. Chalmers 2010, p. 244). This is considered unnecessarily demanding and not to reflect the standards typically accepted in science, given that it does not allow for the possibility of theory correction (Churchland 1989; Schaffner 1993; Bickle 1998). Instead, Hutto and Myin (2013) highlight Churchland’s ‘New Wave’ criteria for reduction, noting that a reduction should rather preserve the ‘*image* of the higher-order theory’ in the lower-order theory on offer, and thus its

⁷ Some enactivists may immediately balk at the sight of (A), most obviously those expounding avowedly “non-reductive” theories of mind (e.g. Colombetti 2014; Thompson 2007). I address this issue in f13.

⁸ As to the significance of Hutto and Myin’s *endorsement* of such criteria, this will be returned to in the second half of the paper.

explanatory and predictive resources (pp. 175–6; Churchland 1989, pp. 49–53).⁹ The higher-order here is the *reduced* theory, while the lower-order is the *reducing* theory. Now, given that the preservation of a theory's explanatory and predictive resources might reasonably be thought to be entailed by the preservation of its *image*, I will focus purely on this latter criterion here. What then does the preservation of an *image* entail?

To make things simpler (and for consistency with the characterisation of reduction given at the beginning of this section) I address this question from the perspective of *ontological reduction*, rather than theory reduction. This is to target the reduction of individual objects or properties referenced within particular theories, rather than whole theories themselves. There is no significant difference in the reductive criteria here; one aims to preserve the *image* of these, in a lower-order explanation. Let us flesh out how this might happen then.

When reducing some higher-level object, a reductive theory will first need to invoke its own specific (i) ontology and (ii) dynamics, i.e. it will offer (i) some lower-level objects and (ii) some account of how these things are disposed to behave (that is to say, it affords them certain properties). It will then cook up an account of the relationship between those objects which can be seen to yield some *macroscopic properties* that mirror the properties of the higher-level entity being targeted for explanation. A macroscopic property here refers to some property instantiated at a scale larger than that of the lower-level objects themselves. Put differently, it is a property that occurs at the *system-level* rather than the component-level. (If it spans the *whole* system, it is called a *global property*). Let us then take an example of the above procedure.

When reducing human biology to cell biology, one might take a human heart (a specific higher-level object, with properties disposing it to behave a certain way). One might then, from within the lower-level theory, outline how assemblies of muscle cells are disposed to coordinate over time via electrical signals so as to create periodically fluctuating global pressure levels across the assembly, a disposition which will itself be seen to mirror the heart's disposition to expand and contract (or beat). Here one is thus invoking some lower-level ontology and dynamics, i.e. some set of cells which are disposed to behave in particular ways, and concocting an account of their relationship that can be seen to yield some *macroscopic property* that is a *mirror image* of higher-level properties of the human heart.¹⁰

We can see that the production of such an image is equivalent to the demonstration of a strong *structural resemblance* between the two relata, for a mirror image can be subject to a relatively detailed abstract level of structural description that is equally applicable to that which it mirrors.¹¹ It is for this reason that I claim SRC is a domain

⁹ Referring to the higher-order theory as the 'older theory', and the lower-order as the 'newer theory', Churchland states: 'A reduction consists in the deduction, within T_N [the newer theory], not of T_O [the older theory] itself, but rather of a roughly equipotent *image* of T_O , an image still expressed in vocabulary proper to T_N [...] the older theory, accordingly, is never deduced; it is just the target of relevantly adequate *mimicry*.' (1989, p. 49).

¹⁰ I am here narrating only how *intrinsic features/properties* of the higher-level can be mirrored. Churchland himself talks of mirroring both intrinsic features/properties and roles/functions within a larger system (1989, p. 52). I do not focus on these latter two here, given that replicating something's intrinsic features will also enable it to play the same causal/functional role (see Drayson 2012, p. 2).

¹¹ Indeed, for something to count as a mirror, it must replicate a fairly *large set* of structural features. The less it preserves, the more *distorting* the mirror and the nearer it approaches a *loss of mirror status*, and thus its capacity to present an *image*.

specific application of this criterion for reduction. If we want, we can thus substitute the vocabulary of reduction into SRC and say that (a) the combined constituent parts and relations invoked by our cognitive science, should comprise a system that instantiates *macroscopic properties* displaying an imagistic relation to the properties invoked in relatum (b) of SRC.

Returning to relatum (b) here, it is important to remember that the target to be mirrored by our cognitive science is always some *conscious process*. This means firstly that the relationship between the constituent parts invoked will need to be a *dynamic* one (given its being a *process*, i.e. something which changes over time) and, secondly, that the target will be provided and characterised *as a target* through phenomenological reflection (given its being a *conscious process*). It is on account of this second and more significant point that SRC invokes *phenomenological characterisations* in its outline; phenomenological theory is the higher-level theoretical domain which constructs the targets that cognitive science is to mirror.¹²

Now that we have all this on the table, we can ask: what is the benefit of SRC? Here we can reintroduce points (A) and (B) from above:

- (A) SRC mirrors a specific criterion for ontological reduction in the sciences that serves to increase the *intelligibility* of making identity claims.¹³
- (B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes.

Point (A) is fairly simple; the production of an *image* of the higher-level, within the conceptually richer lower-level theory, motivates us to believe that both areas of inquiry have indeed managed to latch onto the very same worldly entity (c.f. Churchland 1989, p. 52; Metzinger 2000, pp. 67–69). Put differently, the structural resemblance that has been revealed gives us a strong reason to think that the two levels of description are in fact describing *one and the same thing* in approximately correct ways. Of course, identity relations demand a *complete isomorphism* between two relata. But given that we ought not expect our science to quickly arrive at some “complete and final theory” of the mental, nor ought we to expect even our best phenomenology to be completely immune from any error, a strong structural similarity still provides strong reasons to claim identity (c.f. Churchland 1989, p. 52; Metzinger 2000). Similarly, it seems that the stronger the structural resemblance revealed, the more intelligible it is to make such claims. Strong resemblance relations give a good vertical explanation what Gillett calls ‘ontologically unifying power’ in this way (2016, pp. 225–7). What precisely ought

¹² It will be noted that I have outlined/compared only the detail for a *retentive reduction* in this section, where that which is reduced *retains its conceptual validity*, as when one reduces the biological to the chemical (see Bickle 1992). I do this on account of our explanatory target in SRC being the output of our *best* phenomenology, which I here assume to be accurate enough to resist any eventual elimination. A retentive reduction is thus set in contrast to a *replacement reduction*, which *eliminates* that which it reduces, e.g. in the case of phlogiston theory. I will say more on this important issue of phenomenological accuracy in sections 5, 6 and 8.

¹³ It is worth emphasising here that I am only suggesting that SRC mirrors *one criterion* for reduction in point (A), i.e. *that which I believe to be critically important for motivating identity claims*. One might reasonably argue that this is a necessary but *insufficient* condition for reduction, and that there are additional reasons that the enactivist will wish to construe their approach as non-reductive. See, for example, Thompson (2007) for reasons in favour of a non-reductive construal of enactivism, and Bayne (2004, pp. 258–9) for a contrasting appraisal.

one's identity claims to feature? The set of constituent parts and relations themselves, or the macroscopic features that these display? The answer to this will depend upon one's position on multiple realization. Opting for the former would be to equivocate between a relation of constitution and identity, and to deny any metaphysical difference between the whole and its appropriately arranged parts. Opting for the latter and identifying the mental with (physically-realized) *macroscopic* processes would be to make room for multiple realizability (c.f. Papineau 1998, pp. 377–8). I leave this choice to others, given that it has no bearing on the value of SRC as such.

From (A), we can arrive at the second and more pointed advantage of SRC: (B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes. Specifically, SRC respects the intuition that our subpersonal scientific accounts of the mental and our phenomenological accounts should in some sense *fit together*. This is a criterion that McDowell (1994) has previously offered, recently emphasised by Wheeler (2005, 2013). McDowell (1994) notes that our scientific accounts should not be 'phenomenologically off-key'; they ought not seek to identify the mental with things that do not fit onto (or *match the pitch of*, to extend the musical metaphor) our best phenomenological accounts. What then would it mean to be phenomenologically *on-key*? I think we can see that this is precisely what one is doing when seeking to satisfy Churchland's criteria, which I have shown SRC to be a domain specific version of. Specifically, it is by so constructing an *image* or 'analogue' (Churchland 1989, p. 52) of that which is postulated at the higher (phenomenological) level, *within* the lower (scientific) level, that we match its pitch. Such an imagistic/ analogic relation between the two ensures a phenomenologically *on-key* account.¹⁴ Without such a relation, our cognitive science remains phenomenologically impoverished and we are left feeling that something is missing.

Having outlined SRC then, we are now in a position to return to the questions raised at the end of section 1: (1) Why ought we to think that phenomenological considerations should be able to influence our cognitive science? And (2) in precisely what manner should this occur? We can answer (1) by saying that a scientific account can better motivate identity claims by *doing justice* to the phenomenology (and thus being constrained by it). We can answer (2) by saying that cognitive science can *do justice* to the phenomenology by demonstrating a strong structural resemblance to that which phenomenological reflection sets up.

Let us now return to the issue of perception then, as we began, and apply SRC here. Doing this, we see that we are immediately faced with an important question: what *is* the appropriate phenomenological picture of perception to be mirrored? SRC makes the

¹⁴ The distinction between these two roughly matches McDowell's distinction between a *constitutive understanding* and an *enabling understanding*. A constitutive understanding, Wheeler (2013) notes, 'concerns the identification, articulation and clarification of the conditions that determine what it is for a phenomenon to be the phenomenon that it is' (pp. 142–3). Rather confusingly here, a McDowellian constitutive understanding is thus a case of clarifying what the target for explanation actually is, the domain of phenomenology in the current context. An *enabling understanding* concerns 'the causal elements, along with the organization of, and the systematic causal interactions between, those elements, that together make intelligible to us how a phenomenon of the kind could be realized or generated in a world like ours' (p. 143); this is the domain of cognitive science here. A good enabling understanding can therefore make the constituted object intelligible by offering an account of causal elements and relations that can be seen to realize (rather than generate, in this context) *analogic* macroscopic processes.

answer to this critical. So let us see the enactivist response. Doing so will enable us to understand precisely how the enactive account of perception is derived.

5 Enactivist phenomenology and its consequences

Central to enactivism is the claim that, *phenomenologically*, perceptual experience is an episode of real-time interaction with the environment (Noë 2004; Thompson and Cosmelli 2011; Ward 2012). Interactivity is purportedly invoked in our most rigorous and accurate phenomenological accounts, themselves held to emerge from the methods of (amongst other things) continental phenomenology and Buddhist mindfulness practice (Colombetti 2014; Depraz et al. 2003; Varela and Shear 1999; Varela et al. 1991/1993). Of course, enactivists list numerous additional phenomenological properties of perception—for example, that one half of such interaction comprises a *felt-body*, that acts as a subjective and affectively charged pole of such interaction (see Thompson 2007; Colombetti 2014)—but it is this broad phenomenological structure of *interaction with an environment* (i.e. with something *external* to, or transcending, the subject of perception) that I wish to focus on to serve the purposes of the present argument. Let us entertain some example descriptions then, where this picture emerges.

O'Regan and Noë (2001) exemplify the enactivist stance using the 'feel of a Porsche', an experience that they say is naturally expressed, *phenomenologically*, in terms of the particular 'Porsche-like give and take' between us, the car and the environment (pp. 79–80). When asked to describe such an experience, we talk of Porsche-specific responses to our actions. We talk of *the way the Porsche handles*, the fact that it accelerates quickly when we press on the accelerator for example, pulling us swiftly away with it, and more general facts about both (i) *our own activity* and (ii) *the way the Porsche responds to us*. What we have here is *synchronic interaction*, between subject and object.¹⁵

Let's take another more concrete example: the feeling of *hardness* associated with perceiving a table by touch. This might be described by the enactivist as a sense of *resistance to one's action*. No matter how hard one pushes against the table, the table pushes back. It does not give way. The feeling of hardness should be described in terms of the refusal of that object to be perturbed by the subject, and its exertion of force back upon the subject in response to attempts to distort it.

Enactivists would also wish to extend this kind of phenomenological account to more difficult perceptual experiences, *the apprehension of colour* for example. An enactivist might here wish to highlight the artist Kandinsky's description of the phenomenological difference between yellow and blue: 'if two circles are drawn painted respectively yellow and blue, brief concentration will reveal in the yellow a spreading movement out from the centre, and a noticeable approach to the spectator.

¹⁵ It's important to emphasise that this description concerns the *feel of a Porsche*, not the *feel of the experience*. It has been common in the consciousness literature to sometimes speak of experiences themselves having "feels", or "feeling a certain way to a subject", creating additional complications in explanation (see Hacker 2002; Janzen 2008, chpt. 2). I am generally suspicious of such reference to second-order feels, however, and take all phenomenological descriptions in this paper to be direct accounts of the experiences themselves (i.e. of their qualities).

The blue, on the other hand, moves in upon itself, like a snail retreating into its shell, and draws away from the spectator.’ (1914/1977, pp. 36–7).

Certainly, one might dispute such phenomenological claims—which is an important point we shall return to later—though stating them overtly here should make crystal-clear the further conclusion that enactivists draw about perception.¹⁶ Given enactivists’ widespread commitment to both (i) SRC and (ii) interactive phenomenology, an interactively structured cognitive-science account fits better than a non-interactive one. We have no more than an inference to the best explanation (Rupert 2015, p. 157).

Of course, shared interactivity alone gives only a fairly weak structural resemblance without detail regarding the substructures contained within each pole of interaction, or the kind of interaction involved.¹⁷ Convincing accounts of perception should aim to mirror as many phenomenological features as possible in order to strengthen the resemblance. However, enactivists think that their explanations can also account for just as broad a swathe of these additional features as can competing accounts of perception (which tend to miss interactivity), thus making the resemblance that enactive scientific accounts display stronger, and their explanations more satisfying.

Now that we have these central enactivist commitments on the table—a commitment to (i) SRC and (ii) interactive phenomenology—we can better understand why enactivists make the claims they do about perception. For example, Thompson and Cosmelli remark that ‘given this [interactive] conception of perceptual experience, we can’t specify the mechanisms of perception only in terms of what goes on in the brain without including the body and its dynamic sensorimotor coupling with the environment’ (2011, p. 165). This claim can be understood by noting that it rests on an implicit endorsement of SRC. The authors see a cognitive-scientific explanation yielding global processes labelled in terms of “dynamic sensorimotor coupling” as the most viable means of *mirroring* the interactivity of the phenomenology.¹⁸

Taking SRC as an implicit commitment also enables us to understand similar claims made by Ward (2012, p. 734), who notes that enactivists’ phenomenological commitments make it ‘natural to adopt a *complementary* conception at the subpersonal level [...] it is natural for them to resist a subpersonal conception restricted to the neural activity of the organism engaged in such interaction [for example].’ I suggest that the “natural” pull of such manoeuvres stems from a tacit commitment to SRC, which I

¹⁶ In fact, we shall see that it is crucial for enactivists to hold fast to the principle that phenomenological descriptions *can* be open to error, for such a proposal will undergird their strategies for addressing the explanatory gap (see sections 6 and 8). I am in general agreement with such a position on fallibility. However, it should be noted that, while some stress the pervasive unreliability of introspective judgement across a broad swathe of areas (e.g. Schwitzgebel 2008), the arguments in the remainder of this paper rely mostly on only a narrow range being held questionable. We might want to remain cautious, for example, about questioning the authority of a subject’s judgements concerning the *occurrence* of particular mental states. Yet, I think that we are certainly justified in holding fallible judgements about *that in virtue of which* such a state has a particular content/character (i.e. the structural qualities that, for instance, *make* an experience an experience of hardness). And fallibility concerning these latter types of judgement is the important thing here.

¹⁷ I here, and henceforward, use “interaction” as shorthand for temporally-extended, synchronic interaction between an organism and aspects of the environment that are external to the subject of perception.

¹⁸ A similar example of such an approach to explanation, which talks in terms reminiscent of dynamic sensorimotor coupling, is Bruineberg and Rietveld’s ‘Radical Embodied Cognitive Science’ (2014). This seeks to utilise existing cognitive scientific concepts, such as “affordances”, “sensory feedback”, “prediction error” and “free-energy” etc., in order to offer explanations yielding macroscopic processes labelled as “self-organizing brain-body-environment systems”, mirroring the phenomenology at issue.

have shown to be supported by standalone considerations from the philosophy of science.¹⁹

When entertaining Ward's remark, it is also important to appreciate that combining SRC with an enactivist phenomenology allows enactivism to remain fully compatible with the continued reference to *neural activity* or *neural representations* in cognitive scientific accounts of perception. The point is merely that an enactivist will not seek to *identify* perception with such things. Any cognitive scientific account will always invoke manifold component parts *involved* in the realization of some mental process, which needn't do justice to the phenomenology, and which can be studied as important standalone phenomena (Ward 2012, p. 733). It is perfectly possible therefore for neural representations, for example, to play a role (i.e. at the component level) in the realization of macroscopic interactivity. This is an important nuance recently stressed by Wheeler (2013).²⁰

Now that the nature of SRC has been delineated then, along with its warrant and implications, I wish to approach more critical points. I am concerned specifically with the 'Radical Enactivism' of Hutto and Myin (2013) which, despite seeming to implicitly endorse SRC, makes other claims that I think withdraw from its full implications and threaten to undermine the enactivist project. The problematic claim here is the supposed compatibility of Radical Enactivism with an entirely *narrow* (i.e. brain-bound) supervenience base for perceptual experience. To draw out this problem in the next section, I will first sketch the nature of Hutto and Myin's endorsement of SRC and will then outline the problematic claims. This will enable me to bring out the broader point that enactivists must resist such a manoeuvre, if they wish to maintain that enactivism can address the explanatory gap.

6 Radical enactivism and SRC

Hutto and Myin (2013) have recently set out 'Radical Enactivism'. This is a form of enactivism which *begins with the science*, rather than the phenomenology, and applies considerations about the individuation of cognitive systems in order to recommend an interactively structured cognitive science of the mind. For Hutto and Myin, philosophical critique of our cognitive science can demonstrate that (i) cognitive systems *do not always terminate at the skin*, and (ii) cognitive science need not always utilise the notion of internal 'content' (i.e. representations that are subject to accuracy conditions) in order to satisfactorily explain the mind.²¹ Radical Enactivism therefore falls broadly into line with the central enactivist emphasis upon the need to understand and conceive of the mind in terms of world-involving interaction (p. 4). In this section, I wish to look at some of the claims that Hutto and Myin make about the philosophical implications of such ideas. I focus specifically upon the conclusion of their radical enactivist manifesto, for it is here that we can see their implicit endorsement of SRC.

¹⁹ I shall return to Ward's own reasons for this conclusion in section 7, when they will become important for appraising enactivism's capacity to address the explanatory gap.

²⁰ Though see Gallagher and Zahavi (2008, pp. 93–4) for more sceptical thoughts on the plausibility of 'representation' at this componential level.

²¹ It is the emphatic rejection of the idea the cognition necessarily involves content that is said to afford this version of enactivism its purported radicalism (Hutto and Myin 2013, p. 8).

At the end of their programme for Radical Enactivism, Hutto and Myin come to address the question of the ‘explanatory gap’ (Levine 1983). The explanatory gap is our purported inability to provide or reveal an *intelligible link*, or a *conceptual overlap*, between the physical and the mental, with our explanations (Davies 2008). Hutto and Myin note that any account which fails to do this will leave only *unintuitive identity claims* (2013, chp. 8). The authors also note that they too are out to endorse a strict identity thesis and suggest that ‘belief in such identities can and should be motivated’ (p. 157), thus expressing the belief that the explanatory gap can be closed. To do this, they suggest that most philosophers must go through a process of ‘reconceiving’ (or, more appropriately, *fixing*) the phenomenological accounts that feature in their identity claims (p. 176). They state:

Like other enactivists, we reject the standard ways of characterising the ‘phenomeno’ side of phenomeno-physical identity claims. The difficulty with other existing conceptions of phenomenal properties is that their advocates are wedded to confused pictures of what is to be identified, when they imagine the relata to be qualia and brain states (p. 176)

‘Qualia’ is a notoriously complex term, denoting a particular conception of phenomenal properties that usually ascribes to them at least some of the following set of (second-order) properties: *intrinsically qualitative*, *private*, *ineffable*, *incorrigible*, *atomic* etc. (p. 156; c.f. Dennett 1991/1993). It is unnecessary to go through the entire list here though. What the authors seem to be particularly concerned with in this context is the suggestion that phenomenal properties are both *ineffable* and *non-dynamic*. ‘As long as this picture [of the phenomenal] remains in play,’ they note, ‘there can be no progress in understanding how phenomenality intelligibly relates to, or might be instantiated in, nature’ (p. 157). They suggest instead that we ought to invoke phenomenological accounts that are ‘in tune with’ (ibid.) the way that we naturally speak about our experiences. They remark, ‘[n]aturally occurring ‘what it feels like’ illocutions take activities as their natural objects. When we describe phenomenal experience, we cannot help but mention environment-involving interactions’ (p. 177). As such, we see the same phenomenological picture recommended as was outlined in section 5. What’s more, it is important to stress that we are asked to reconceive our phenomenology in such an interactive manner *because these accounts are supposedly more accurate*, not simply because they are natural, nor because they help us argumentatively. But what argumentative function does this “revised” (or fixed) phenomenology serve? How does it help us work against the explanatory gap?

Hutto and Myin continue, in what is I think is the crucial comment, stating that ‘the plausibility of the proposed identities [been the physical and the phenomenal] looks entirely different, and far less contrived, if it is assumed that the phenomenal character of experiences must, ultimately, be understood by appealing to interactions between experiencers and aspects of their environment’ (2013, pp. 176–177). This seems to me to be nothing other than a commitment to SRC; we can lessen contrivances, and therefore motivate identity claims, by making our relata appropriately resemblant in structure. Having recommended an interactively shaped cognitive scientific account of perceptual experience, they suggest an interactively shaped phenomenological one with which to pair it. Hutto and Myin remark that they ‘foreground the ways in which

environment-involving activities are required for understanding and conceiving of phenomenality [and] abandon attempts to explain phenomeno-physical identities in deductive terms for attempts to motivate belief in such identities by reminding us of our common ways of thinking and talking about phenomenal experience.’ (2013, p. 177) Here, we therefore see an attempt to provide conceptual overlap between the mental and physical via a revisionary stance on phenomenal properties.

We can also see here that Hutto and Myin are therefore abiding by SRC from the *reverse direction* to Thompson and Cosmelli (2011). Having spent their outline of Radical Enactivism arguing for the need to approach the mind *subpersonally* (in cognitive science) in terms of world-involving interaction—using metaphysical arguments about the individuation of systems and what they call the ‘Hard Problem of Content’ (p. xv)—they *then* invoke interactive phenomenology *post-hoc* as a means to generate structural resemblance and validate such a subpersonal framing of perception.²² Yet in so doing, they seem to be promoting SRC in essentially the same manner. We can construe this kind of argumentative manoeuvre as a direct counter to the concerns of Rupert (2015), for example. Evaluating the strength of enactivist theories, Rupert draws an unfavourable assessment of their supposed progress against the explanatory gap through invoking subpersonal interactions, asking ‘[h]ow could that stuff—interaction with the environment, for example—*be* phenomenal experience? It doesn’t *seem* at all like consciousness’ (2015, p. 161).²³ Hutto and Myin would be apt to respond here by saying ‘yes, it doesn’t *seem* like experience, because you are working with a confused phenomenology’. Rupert is unpersuaded because he seems to take the phenomenal properties of perception to be non-dynamic (i.e. non-interactive) while, for the enactivist, dynamism is *characteristic* - it’s dynamism all the way down. Rupert thus perceives an explanatory gap to remain because of his conception of the phenomenal. Radical Enactivism seeks to work against the explanatory gap by offering a *re-conception* of the phenomenal.²⁴ So far, so good.²⁵

Yet all is not well with Radical Enactivism. My concern here is Hutto and Myin’s insistence that Radical Enactivism is compatible with the *narrow supervenience* of perceptual experience. Examining this now can help bring out a larger point about the

²² Clearly, Hutto and Myin are invoking a different kind of phenomenological source to that of Thompson, for example, who invokes the results of more disciplined kinds of investigation (e.g. 2007, 2015; Thompson and Cosmelli 2011; Varela et al. 1991/1993). See section 8 for more on this issue.

²³ Despite being unclear in this quote, the enactivist need not construe all forms of experience in terms of interaction. I am concerned here only with perception, which is argued to have such character, but other experiences (e.g. pain, or emotion) need not.

²⁴ In this sense then, Radical Enactivism joins other enactivists in rejecting any *metaphysical* gap between the mental and the physical (see e.g. Noë and O’Regan 2002). It maintains the existence only of an *epistemic* gap, occurring at the level of individuals, which can be appropriately addressed (i.e. closed) by providing those individuals with better phenomenology.

²⁵ Rupert has a further counter to enactivist suggestions here, arguing that structure is tangential to the issue at hand. Linking the explanatory gap to the hard problem (given that bridging the explanatory gap would solve the hard problem) he states, ‘the hard problem is that of explaining *the nature of things on which structure is imposed, not the structure itself*’ (2015, p. 161, emphasis added). As such, given that enactivist phenomenology only speaks in terms of structured dynamics, they can only leave the gap as large as ever. Bayne (2004, p. 361) also posits this same predicament. In response, we might ask: what relevant things *could* we say about the ‘nature’ of conscious processes beyond their structure and causal efficacy? It is unclear what is being demanded here. It seems equivalent to asking a scientist to explain the intrinsic nature of some physical entity beyond its compositional structure and causal powers. Perhaps there is something else important here, but I do not see it.

commitments enactivists must make in order to retain the full philosophical force of their ideas.

7 Radical enactivism: correcting inconsistencies

When assessing the explanatory benefits of enactivism's proposal to "go wide" with perceptual experience— i.e. to appeal to world-involving processes of interaction in subpersonal (cognitive scientific) explanation—Hutto and Myin state that '[t]he explanatory value of this move is not undermined, even if it is accepted that the supervenience basis of phenomenality is wholly brain-bound' (p. xix).²⁶ Though they do not endorse the narrow-supervenience thesis, they are keen to suggest that their Radical Enactivism is *compatible* with it. This kind of move goes directly counter to positions held by many other contemporary enactivists.

A majority of prominent enactivists support the Extended Conscious Mind Thesis (ECM): the material (constitutive) supervenience base of some kinds of experience extends into the environment (Hurley 1998; Noë 2004; Pepper 2014; Thompson and Cosmelli 2011; Ward 2012). The typical candidate here is *perceptual experience* (Pepper 2014, p. 99). ECM about perceptual experience is thus a commitment to a *wide* supervenience base.²⁷

Hutto and Myin wish to retreat from this. Instead, they offer a metaphysical thesis restricted to 'basic cognition'. Basic cognition refers to foundational forms of cognition, including such things as cricket phonotaxis (pp. 42–43), but also relatively simple aspects of human cognition. It is defined as 'mental activity that exhibits intentional-directedness, but it doesn't necessarily imply phenomenality [i.e. phenomenal properties]' (p. x). They claim that such basic cognition is 'extensive', i.e. 'fundamentally, constitutively already world-involving' (p. 137). Yet they stress that this does not entail a commitment on their part to *conscious* mental activity (such as perception) being itself extensive.

This theoretical retreat seems highly problematic to me. I think that there is good reason that enactivists tend towards ECM, for I suspect that the 'explanatory value' of enactivism is indeed put in danger by Hutto and Myin's stance here. What I find problematic is not so much Hutto and Myin's claim about basic cognition, it is *the claim that they can avoid ECM* (a thesis of wide supervenience), in conjunction with other claims they have made about enactivism's explanatory value. As such, we must first understand what they mean by 'explanatory value'. This begins to be revealed in the below quote:

It remains an open question whether the strategy of "going wide" for explanatory purposes implies that the minimal supervenience base for phenomenality is like-

²⁶ Though this quote is ambiguous, an interpretation of it as thematising *perceptual* phenomenality/experience, rather than only *some* (other) kinds of (non-perceptual) phenomenality/experience, is motivated by its later evaluation in the context of 'ordinary on-line perceiving' (Hutto and Myin 2013, p. 161).

²⁷ Note that on the current rendering a rejection of narrow supervenience does not entail a commitment to ECM. ECM concerns the environment, while narrow supervenience concerns only the brain. One could argue that perceptual experience supervenes on brain and (non-neural) body, but not the environment. However, given that Hutto and Myin's phenomenology is interactive, it is ECM that I focus on here.

wise (i.e. extensive). Although *a full and satisfying understanding* of phenomenality cannot be achieved without going wide, it is compatible with the supervenience base for phenomenality remaining entirely confined to the brain (2013, pp. 157-8, emphasis added)

We see here that it is *understanding* that is at stake. Going wide provides full and satisfying understanding; this is its explanatory value. Thus, enactivism's explanatory value is preserved if it can offer such a full and satisfying account. However, it seems that the authors' attempts to make Radical Enactivism compatible with narrow supervenience invoke *two different kinds of understanding*, which are themselves inconsistent. Only one of these is full and satisfying. Furthermore, this kind is *incompatible* with narrow supervenience. My suggestion is therefore that Hutto and Myin cannot claim to provide both of these forms of understanding; they must choose between them based upon the philosophical work that they want Radical Enactivism to do. Let me unpack this.

On the one hand, Hutto and Myin suggest that "going wide" provides a *contextual* understanding. They note that 'the strategy of going wide is necessary when trying to understand the phenomenality of experience *in a wider context*' (p. 165, emphasis added). This is to say, such a cognitive-scientific story can tell us about the specific sensorimotor patterns of interaction between organisms and environments that are *ordinarily needed in place* for such experiences to occur (Ibid.), and which perceptual experience has been historically tied to. This is to suggest that "going wide" with our cognitive science can give us the typical *contextual conditions* of perceptual experience, whether these be real-time causal conditions, or historical conditions, without taking these to be constitutive of perception, or 'part of its metaphysical essence' (Ibid.). Hutto and Myin rightly note that such contextual understanding is useful on many counts, especially if we wish to understand how such experience arises in nature, or the conditions under which it can be most easily generated (pp. 163–4).

This first kind of understanding is invoked while Hutto and Myin entertain challenges to ECM. They believe that the most powerful way to problematise ECM is the 'Argument from Shared Phenomenality' according to which very similar, if not identical experience can occur in both the presence *and the absence* of organism-environment interaction. For example, we might consider extravagant Swampman thought experiments, where the neural activity of a perceiver is brought into being from nowhere (Davidson 1987), or 'direct neural manipulations of envatted brains' (Hutto and Myin 2013, pp. 161–2). The suggestion here is that we can coherently imagine the same kind of experience being preserved in the absence of real-time environmental interaction, with appropriate metaphysical conclusions then being drawn.

Alternatively, the Argument from Shared Phenomenality can be made by referencing more mundane cases of supposed experiential overlap in the apparent absence of environmental interaction, such as 'individuals dreaming, imagining, and suffering from "locked in syndrome"' (ibid). We might also imagine such things as subjectively indistinguishable hallucinations. Broadly, all such arguments rest upon motivating the possibility of shared phenomenality across cases of both online perceiving, where real-time environmental interaction is in place, and (often less epistemically 'good') cases where that interaction is absent. Hutto and Myin are suspicious of such

arguments, yet want Radical Enactivism to remain standing even if these go through (pp. 161–3). They thus use the hypothetical soundness of such arguments to motivate invoking the above kind of *contextual understanding*.

On the other hand, however, we have a second form of understanding offered by Hutto and Myin: a *constitutive understanding*. We have already seen that the authors are keen to motivate identity claims between the mental and the physical (p. 157), and suggest that this can be done by bringing our cognitive science and phenomenology into a relation of structural resemblance. However, this can only motivate identity claims if we take the science to be giving us a *constitutive account* of a mental process, rather than a contextual account of how it came to be. Returning to an earlier distinction, this means we must be offering a *vertical* rather than a *horizontal* explanation. As such, Hutto and Myin here tie the explanatory value of enactive cognitive science directly to its offering (i) a constitutive, vertical account that (ii) *abides by SRC*. In this case then, they cannot divorce the explanatory value of enactivism from metaphysical claims about relations of supervenience (i.e. *constitutive supervenience*, here).

Additionally, I have shown SRC to be a domain-specific application of Churchland's criterion for reduction. And this is a criterion which seems to satisfy Hutto and Myin. Indeed, they themselves promote this conception of reduction as a means to motivate identity claims, stating that we should bring our explanatory criteria 'in line with this more liberal thinking about how identities can be established' (p. 176), further motivating the thought that this is indeed the kind of understanding they are keen to provide, and again suggesting a supervenience entailment.

Having got these two kinds of understanding on the table, it is first immediately clear that they are incompatible. *Either* one thinks that "going wide" with our cognitive science is telling us about the causal/historical context of perceptual experience *or* one thinks that it is telling us about its constitution and thus its supervenience base; one cannot claim both from one explanation. So we have a curious oscillation here, on Hutto and Myin's part, between two different and incompatible forms of understanding.²⁸ Which one ought they then to favour? It seems clear to me that, given Hutto and Myin's claims about the explanatory gap, they ought to favour dropping the claims about contextual understanding. Why so? Because it is only within the realm of a constitutive understanding that we can hope to bridge the explanatory gap. Tackling the gap requires accounting for *what the mental is* in terms of the physical, not accounting for how it came about. It requires demonstrating conceptual overlap between things that are *already proposed as candidates for identity*. And this is indeed what Radical Enactivism sometimes seems to aim at in its demonstration of structural resemblance between specific "interactive" phenomenological construals of perception, and cognitive-scientific renderings of these same things that manifest an analogic interactive structure at the global-level. What's more, it seems to be in this engagement with the explanatory gap that the explanatory value of enactivism is properly secured. Hutto and Myin want to offer explanations that provide "full and satisfying" understandings. If they do not address the explanatory gap, they do not do this.

A second point can be made against the idea that enactivist cognitive science should be construed as offering a merely contextual account. By definition, a contextual

²⁸ See Froese (2014) for a similar assessment.

account will invoke interactive subpersonal processes that are *not themselves constitutive* of some mental process. A proposed contextual account of this kind would then necessarily admit the possibility of some *additional* constitutive subpersonal story about perceptual experience that *could not feature subject-object interaction* (given that interaction is meant to *frame* this, and cannot therefore figure *within* it). Yet, given Hutto and Myin's emphasis upon an interactive phenomenology, this would dissolve the central resemblance that they elsewhere invoke in support of their ideas (c.f. Ward 2012, pp. 734–5). So not only does their promotion of this kind of contextual understanding fail to bridge the explanatory gap, it actively works against the way they elsewhere suggest closing it.

This allows me to make a broader point than about enactivism in general. If enactivists wish to maintain that they *can* address the explanatory gap, they simply cannot be concessionary in the way that Hutto and Myin suggest. If enactivism is allowed to be compatible with narrow supervenience, *in spite of insistence upon an interactively-structured phenomenology of perception*, it will need to drop any endorsement of SRC, and be content to offer more limited kinds of understanding.²⁹ This, however, would be to disarm enactivism of its major weapon against the explanatory gap. Perhaps this is an appealing strategy for some, though I cannot imagine it would satisfy Hutto and Myin, given their otherwise valuable remarks on this issue.³⁰

A final way to motivate the above choice amongst the broader enactivist community (opting for a *constitutive* construal of enactive cognitive science) is to note that it helps enactivists retain their contention that phenomenology can have some bite. We can stress here that there would be no reason at all to endorse SRC, as a means of addressing the explanatory gap, unless one believed that one's phenomenological accounts were themselves *accurate* in some sense, that is to say, unless one believed that they truthfully revealed *something* about the subpersonal make-up of perception. This point is made by Noë (2007) who states that, if one wishes to avoid the 'epistemic isolation' of phenomenology, one must be prepared to say that phenomenology makes a *theoretical commitment*, which is accordingly subject to pressure from science (p. 232). After making the phenomenological claim that perceptual experience is 'world-involving', he states that 'if it turns out that it is possible for me to have an experience of the same kind as a visual experience of a pencil in the absence of a pencil, then it turns out that I am wrong about my phenomenology itself [...] [for] I take a stand on the theoretical question when I take a stand on the phenomenology' (2007, p. 236).

²⁹ Of course, opting for this more ambitious kind of understanding means that enactivists will need to offer ways of countering the Argument from Shared Phenomenality. They might take, for example, the same line as disjunctivists on this issue, denying that veridical perception and illusion/hallucination are of a common kind. See also Thompson and Cosmelli (2011) and Hutto and Myin (2013, pp. 161–162) for further suggestions on this issue.

³⁰ There is an additional argument against SRC which may further motivate some enactivists to drop it. This is due to Wheeler (2015) who notes that a constraint along SRC lines contradicts the occasional assertions of sensorimotor enactivists that real-time interaction is not strictly *necessary* for perceptual experience. Instead, the thought goes that one only needs a practical mastery of the laws of sensorimotor contingencies that *would* obtain were one to engage in certain motor behaviours (i.e. the possession of certain practical *sensorimotor knowledge*) (pp. 169–173; cf. O'Regan 2011). Given Hutto and Myin's deeply sceptical stance on the minimal necessity of such knowledge however (2013, pp. 24–32), I think they are unlikely to find such an argument convincing. See also Ward (2012, f7) for a response to Wheeler along these same lines.

For these reasons, Ward (2012) argues that once enactivists have an interactive phenomenology in play this makes it natural to resist any narrow subpersonal construal of perceptual experience (and in turn favour ECM), for ‘the picture of the subpersonal that would result from doing so threatens to make it unintelligible how their conception of things at the personal level could be correct’ (p. 735). And surely enactivists like Hutto and Myin *do* want to maintain that their proposed phenomenology is correct. If this were not so, it would sit very oddly with their emphasis upon the importance of *reconceiving* our phenomenology. Why would it matter how we conceive of the phenomenology if its descriptions could be so readily bypassed?

Noë uses the above thoughts to highlight what he thinks is the contrast between continental phenomenology as a disciplined examination of lived experience and the “introspective” tradition of phenomenological reflection, which too often conceives it simply as a free-standing exercise - a ‘descriptive preliminary to theorising’ (p. 232) that metaphysics then comes to supersede. And it is occasionally this kind of free-standing conception that Hutto and Myin seem to be in danger of falling into when they back-slide into claims about contextual understanding.

It must thus be emphasised that Noë’s contention that phenomenology makes a theoretical commitment is a *condition* for endorsing SRC. Unless one holds one’s phenomenology to make claims about the nature of the world, one has no reason for thinking that phenomenology should have any influence when it comes to cognitive science. As such, if the enactive project has any place for phenomenology it must, as Noë remarks, conceive phenomenology to be ‘concerned with nature itself’ not only with ‘how things seem’ (p. 234).

This brings to an end the substantive claims I wish to make in this paper. I now wish to conclude on a more cautious note. For though I think that SRC is warranted, and that the way it is wielded by enactivists suggests a promising route forwards, it must be admitted that a number of significant concerns remain to be addressed here.

8 Questions going forwards and neurophenomenology

First and foremost, given that SRC affords phenomenology an ability to constrain cognitive science, one must be careful to ensure it is being constrained by the *right phenomenology*. Recapitulating an earlier worry: who is to say that an interactive phenomenological account of perception is the correct one? One might claim that phenomenological methods reveal perceptual experiences to be no more than mental paint (Block 1990) or unequivocally representational in nature (Chalmers 2010, p. 334). This is an area of high controversy and complex debate, where arguments require the extended justification of phenomenological methods that can work not merely at the broad level, but also at the finer level, to isolate precisely what is central to perceptual experience from what is perhaps tangential or contextual.³¹ The drawing of these lines is no easy task.

Enactivists contend that certain individuals are well equipped to perform this function, which does not in itself seem an unreasonable suggestion. But what are the

³¹ See Bayne and Spener (2010) for a more detailed sketch of this debate as manifest across the study of a wider array of our experiences.

best sources for maximally accurate and rigorous phenomenology? Certainly, suspicions are likely to be raised by Hutto and Myin's invocation of naïve or "natural" introspective judgement as an appropriate source of phenomenological constraint. This is highly unlikely to persuade without a great deal of argument that naïve descriptions of experience somehow gain access to some wellspring of phenomenological truth that abstract philosophising manages to miss.³²

Perhaps the methods of continental phenomenology or Buddhist mindfulness meditation are more appropriate than here. Yet these are subject to their own criticisms. In particular, there are concerns that such methods, as forms of mental training, also *change* a subject's experience, making it problematic to generalise from the phenomenological reports of such persons (e.g. Rupert 2015, p. 172, f9). A number of recent arguments have sought to assuage such worries, suggesting that transformation of the mind through training in no way rules out gaining insight into its pre-transformed character. Colombetti (2014, p. 156) for example, agrees with Gallagher and Zahavi (2008) that such training always 'involves [both] a gain and a loss' (p. 63), bringing out certain aspects of experience, ordinarily implicit and hidden, at the expense of distorting other aspects of that experience. However, further arguments can be offered to explicate the scope of such gains and the significance of such losses, given that these methods are playing such an important role in the theoretical work being produced.

In addition, further material can be offered by enactivists to crystallize precisely *how* we can determine the accuracy of first-person methods. Existing options here include: (i) intersubjective validation (Gallagher and Zahavi 2008) (ii) introspection-reliant abilities (Spener 2015) and (iii) subpersonal constraints (e.g. Varela 1996), where third-personal data is also used as a restraint upon our phenomenology. While I do not wish to arbitrate between these, I will end with a few remarks on the final of this trio of options, for this possibility is particularly important to keep in mind for those who might worry that SRC affords phenomenology an unreasonably strong influence.

It must be emphasised here that SRC does *not* leave our phenomenology immune from constraints working in the opposite direction, namely from the scientific to the phenomenological. In fact, SRC leaves open that our *best* phenomenological characterisations might themselves need to be determined with the help of subpersonal data. I will thus end here with some brief comments on the *neurophenomenological movement*, often allied to enactivism, which has begun to delineate a way of combining SRC with this opposite direction of constraint, resulting in a method of so-called 'mutual constraints' (Varela 1996).

Thompson et al. tell us that '[t]he [first] working hypothesis of neurophenomenology in an experimental context', which should be familiar given the material covered in this paper, 'is that phenomenologically precise first-person data produced by employing first-person/second-person methods provide strong constraints on the analysis and interpretation of the physiological processes relevant to consciousness' (2005, pp. 46–7). In short, neurophenomenology *assumes* the value of SRC and looks for subpersonal physiological dynamics showing structural resemblance to aspects of

³² This is not to say that this is not a potentially fruitful avenue to pursue. Given that such conceptions (folk and philosophical) are neither univocal nor atemporal, sociological and historical investigations of the diverse, shifting understandings of human experience (and indeed "the physical") will likely prove valuable tools for addressing the explanatory gap. This would be to venture into the realm of *experimental philosophy*.

the phenomenology. Doing so, it aims to reveal previously unrecognised physiological dynamics that are critical to the mental process being thematised. In a famous example, Lutz et al. (2002) examined the experience wherein a subject gradually perceives a three-dimensional figure emerging from an autostereogram, showing that subtle transitions between various stages of this experience (described from the first-person) were mirrored by transitions into distinct stages of “phase synchrony” across different regions of the brain.³³

It is particularly important to note here though that neurophenomenology was also described as a method that would employ constraints in the reverse direction, i.e. it would also use *subpersonal data* to help evaluate the phenomenological descriptions on offer and to make amendments to invoked phenomenology (Varela 1996). Thompson et al. (2005) note that this data can ‘provoke revisions and refinements of the phenomenological accounts’ (p. 47). For example, should a phenomenological report suggest structural features that seem strongly discordant with those found in the subpersonal dynamics, the subject might be asked to re-examine this aspect of their experience. Attending more intimately to this part of experience might then yield different phenomenological reports. However, Thompson et al. go on to note that appropriately discordant third-person data can also give the subject a means of latching onto ‘previously inaccessible’ aspects of their experience (ibid.). It can give the subject *something to look for*. When (or perhaps, if) this something is found and examined in more detail from the first-person perspective, this might allow *richer phenomenological details* about this part of the experience to emerge, that can then be used, in turn, to discover (i.e. look for) relevant (structurally-resemblant) finer-grained subpersonal dynamics. Thus, by combining the two directions of influence, we are supposed to use phenomenological and scientific data as ‘reciprocal constraints’ to help us improve each one, and home in upon what is central (Varela 1996, p. 343). Engaging in this neurophenomenological back-and-forth is thus intended to catalyse the production of ever stronger structural resemblances, which Varela himself saw as part of a ‘methodological solution to the hard problem’ (ibid.).

Unfortunately, this reverse constraint from science to phenomenology is seldom evidenced in neurophenomenological experiments, and far more evidence of the efficacy of the method of mutual constraints is needed. Furthermore, even with a method of reciprocal constraints, we have a problem, for it leaves us uncertain precisely *when* we are supposed to reverse the direction of constraint. In other words, how long do we allow our phenomenology to hold sway as a constraint upon our science, instead of applying the constraint in the reverse direction?³⁴ We must remember that

³³ Of course, this experiment demonstrates only a weak resemblance between phenomenology and *neural dynamics*, capturing a shared *temporal structure*. To strengthen structural resemblance between phenomenology and the subpersonal, enactivists believe we will need to extend our constitution base beyond the brain. This is something that Colombetti (2014) notes in her programme for affective science (the study of feeling/emotion): ‘observing brain activity provides only a partial glance into the organismic activities that enact lived experience [...] This is also why neurophenomenology, which has so far limited its third-person methods to brain activity, ought to include recordings of bodily activity—thus becoming what I shall call [...] *neuro-physio-phenomenology*’ (p. 143). And in the case of perceptual experience, if we buy an interactive phenomenology, a constitution base will need to include the organism’s environment, delineated in terms akin to the ‘Radical Embodied Neuroscience’ of Bruineberg and Rietveld (2014).

³⁴ Thanks to Mike Wheeler (personal correspondence) for this point.

resemblance can be generated by modifying *either* the phenomenal or the subpersonal relata, and thus be prepared to do either one.³⁵

This relates to another problem that enactivists more broadly must deal with. Namely, enactivism's preferred method of operationalising SRC—neurophenomenology—seems to construe third-person constraints as acting primarily as *refinements* to the phenomenology, i.e. as suggesting alterations *within* their interactive phenomenological picture, rather than allowing any major overhaul of this interactive conception, thus yielding a rather a one-sided rendering of the mutual constraint method. Perhaps the neurophenomenologist is warranted in replying to the effect that “look, we need to begin somewhere and, given that we're interested in conscious processes, and that these only emerge *as targets* from phenomenological reflection, we should give this a certain priority in characterising the targets.” Nevertheless, such reply leaves open important questions about the precise range of judgements for which phenomenological reflection will have priority, and more must be said about how to non-arbitrarily carve off such a region.

Finally, given any such neurophenomenological commitment to the authority of phenomenology in such domains (i.e. concerning the broad shape of our experiences), it is evident that something distinct from third-personal data will be needed to ascertain the validity of such first-person descriptions. This might concern the methodological criteria for good phenomenological reflection itself, which is something that other neurophenomenology enthusiasts have picked up on recently (see e.g. Bitbol and Petitmengin 2016; Petitmengin 2009; Petitmengin and Bitbol 2009), though this will need fleshing out more fully to satisfy sceptics.

These are all questions that the contemporary enactivist must address if they wish to continue making use of SRC in the manner that they have done so thus far. But however these questions are answered, I believe that the enactivist promotion of SRC as a standalone constraint is a reasonable one. What's more, I hope that SRC's outline here, if it has any take-home message, demonstrates the indispensability of phenomenology to cognitive science and philosophy, and the importance of doing it well.

Acknowledgements I am much indebted to Maja Spener and Scott Sturgeon for their invaluable advice on numerous drafts of this paper. Thanks must also go to those who attended the 2016 Annual Graduate Conference in Phenomenology, ‘Phenomenology, Science and Consciousness: The Transcendental and the Naturalistic’, where an early version of this paper was presented. I am likewise grateful for the comments and suggestions of two anonymous reviewers.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bayne, T. (2004). Closing the gap? Some questions for neurophenomenology. *Phenomenology and the Cognitive Sciences*, 3(4), 349–364.
- Bayne, T., & Spener, M. (2010). Introspective humility. *Philosophical Issues*, 20(1), 1.

³⁵ It seems to be for this reason that Noë asserts that ‘phenomenology doesn't fix the outcome of important theoretical investigations [...] phenomenology bears on such disputes without fixing them, just as other information about how things are may bear.’ (2007, p. 237, emphasis added)

- Bermúdez, J. L. (2006). *Philosophy of psychology: contemporary readings*. New York: Routledge.
- Bickle, J. (1992). Mental anomaly and the new mind-brain reductionism. *Philosophy of Science*, 59, 217–230.
- Bickle, J. (1998). *Psychoneural reduction: the new wave*. Cambridge: MIT Press.
- Bitbol, M., & Petitmengin, C. (2016). On the possibility and reality of introspection. *Mind and Matter*, 14(1), 51–75.
- Block, N. (1990). Inverted Earth. In J. Tomberlin (Ed.), *Philosophical perspectives* (Vol. 4, pp. 53–79). Atascadero: Ridgeview.
- Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8(599), 1–14.
- Chalmers, D. J. (2010). *The character of consciousness*. New York: Oxford University Press.
- Churchland, P. S. (1989). *Neurophilosophy: toward a unified science of the mind-brain*. Cambridge: MIT Press.
- Colombetti, G. (2014). *The feeling body: affective science meets the enactive mind*. Cambridge: MIT Press.
- Cummins, R. C. (1983). *The nature of psychological explanation*. Cambridge: MIT Press.
- Davidson, D. (1987). Knowing one's own mind. *Proceedings of the American Philosophical Association*, 61, 441–458.
- Davies, M. (2008). Consciousness and explanation. In L. Weiskrantz & M. Davies (Eds.), *Frontiers of consciousness* (pp. 1–54). Oxford: Oxford University Press.
- De Jesus, P. (2016). Autopoietic enactivism, phenomenology and the deep continuity between life and mind. *Phenomenology and the Cognitive Sciences*, 15(2), 265–289.
- Degenaar, J., & O'Regan, J. K. (2015). Sensorimotor theory and enactivism. *Topoi*, 1–15. doi:10.1007/s11245-015-9338-z
- Dennett, D. C. (1991/1993). *Consciousness explained*. London: Penguin.
- Depraz, N., Varela, F. J., & Vermersch, P. (2003). *On becoming aware: a pragmatics of experiencing* (Vol. 43). Amsterdam: John Benjamins Publishing.
- Di Paolo, E., Rohde, M., & De Jaegher, H. (2010/2014) Horizons for the enactive mind. In J. Stewart, O. Gapenne & E. Di Paolo (Eds.), *Enaction: Toward a new paradigm for cognitive science* (pp. 33–87). Cambridge: MIT Press.
- Drayson, Z. (2012). The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, 26(1), 1–18.
- Froese, T. (2014). Review of Hutto and Myin “Radicalizing enactivism: basic minds without content”. *The Journal of Mind and Behavior*, 173(4), 466–500.
- Gallagher, S., & Zahavi, D. (2008). *The phenomenological mind*. New York: Routledge.
- Gillett, C. (2016). The metaphysics of science, nature and the rules of engagement. In K. Aizawa & C. Gillett (Eds.), *Scientific composition and metaphysical ground* (pp. 205–248). London: Palgrave MacMillan.
- Hacker, P. (2002). Is there anything it is like to be a bat? *Philosophy*, 77(02), 157–174.
- Hatfield, G. (2005). Introspective evidence in psychology. In P. Achinstein (Ed.), *Scientific evidence: philosophical theories and applications* (pp. 259–286). Baltimore: John Hopkins University Press.
- Hurley, S. L. (1998). *Consciousness in action*. Cambridge: Harvard University Press.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing enactivism: basic minds without content*. Cambridge: MIT Press.
- Janzen, G. (2008). *The reflexive nature of consciousness* (Vol. 72). Amsterdam: John Benjamins Publishing.
- Kandinsky, W. (1914/1977). *Concerning the spiritual in art*. New York: Dover Publications.
- Kriegel, U. (2015). *The varieties of consciousness*. New York: Oxford University Press.
- Levine, J. (1983). Materialism and qualia: the explanatory gap. *Pacific Philosophical Quarterly*, 64(4), 354–361.
- Lutz, A. (2002). Toward a neurophenomenology as an account of generative passages: a first empirical case study. *Phenomenology and the Cognitive Sciences*, 1(2), 133–167.
- Lutz, A., Lachaux, J. P., Martinerie, J., & Varela, F. J. (2002). Guiding the study of brain dynamics by using first-person data: synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the National Academy of Sciences*, 99(3), 1586–1591.
- McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quarterly*, 44(175), 190–205.
- Merleau-Ponty, M. (1945/1962). *Phenomenology of Perception* (trans: Smith, C.). London: Routledge Press.
- Metzinger, T. (2000). *Neural correlates of consciousness: empirical and conceptual questions*. Cambridge: MIT Press.
- Noë, A. (2004). *Action in perception*. Cambridge: MIT Press.
- Noë, A. (2007). The critique of pure phenomenology. *Phenomenology and the Cognitive Sciences*, 6(1–2), 231–245.

- Noë, A., & O'Regan, J. K. (2002). On the brain-basis of visual consciousness: a sensorimotor account. In A. Noë & E. Thompson (Eds.), *Vision and mind: selected readings in the philosophy of perception* (pp. 567–598). Cambridge: MIT Press.
- O'Regan, J. K. (2011). *Why red doesn't sound like a bell: understanding the feel of consciousness*. New York: Oxford University Press.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(05), 939–973.
- Papineau, D. (1998). Mind the gap. *Noûs*, 32(S12), 373–388.
- Pepper, K. (2014). Do sensorimotor dynamics extend the conscious mind? *Adaptive Behavior*, 22(2), 99–108.
- Petitengin, C. (2009). The validity of first-person descriptions as authenticity and coherence. *Journal of Consciousness Studies*, 16(10–1), 252–284.
- Petitengin, C., & Bitbol, M. (2009). Listening from within. *Journal of Consciousness Studies*, 16(10–1), 363–404.
- Rupert, R. D. (2009). *Cognitive systems and the extended mind*. New York: Oxford University Press.
- Rupert, R. D. (2015). Embodiment, consciousness, and neurophenomenology: embodied cognitive science puts the (first) person in its place. *Journal of Consciousness Studies*, 22(3–4), 148–180.
- Schaffner, K. F. (1993). *Discovery and explanation in biology and medicine*. Chicago: University of Chicago Press.
- Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review*, 117(2), 245–273.
- Spener, M. (2015). Calibrating introspection. *Philosophical Issues*, 25(1), 300–321.
- Stewart, J. (2010/2014). Foundational issues in enaction as a paradigm for cognitive science. In J. Stewart, O. Gapenne & E. Di Paolo (Eds.), *Enaction: toward a new paradigm for cognitive science* (pp. 1–31). Cambridge: MIT Press.
- Thompson, E. (2007). *Mind in life: biology, phenomenology, and the sciences of mind*. Cambridge: Harvard University Press.
- Thompson, E. (2015). *Waking, dreaming, being: self and consciousness in neuroscience, meditation, and philosophy*. New York: Columbia University Press.
- Thompson, E., & Cosmelli, D. (2011). Brain in a vat or body in a world?: Brainbound versus enactive views of experience. *Philosophical Topics*, 39(1), 163–180.
- Thompson, E., Lutz, A., & Cosmelli, D. (2005). Neurophenomenology: an introduction for Neurophilosophers. In A. Brook & K. Akins (Eds.), *Cognition and the brain: The philosophy and neuroscience movement* (pp. 40–97). New York: Cambridge University Press.
- Titchener, E. B. (1899). *An outline of psychology*. New York: Macmillan.
- van Riel, R., & Van Gulick, R. (2016). Scientific Reduction. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), URL = <<http://plato.stanford.edu/archives/spr2016/entries/scientific-reduction/>>.
- Varela, F. J. (1996). Neurophenomenology: a methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3(4), 330–349.
- Varela, F. J. (1997). The naturalization of phenomenology as the transcendence of nature: searching for generative mutual constraints. *Alter: Revue de Phenomenologie*, 5, 355–385.
- Varela, F., & Shear, J. (Eds.). (1999). *The view from within: first-person approaches to the study of consciousness*. Bowling Green: Imprint Academic.
- Varela, F. J., Thompson, E., & Rosch, E. (1991/1993). *The embodied mind: cognitive science and human experience*. Cambridge: MIT Press.
- Ward, D. (2012). Enjoying the spread: conscious externalism reconsidered. *Mind*, 121(483), 731–751.
- Wheeler, M. (2005). *Reconstructing the cognitive world: the next step*. Cambridge: MIT press.
- Wheeler, M. (2013). Science friction: phenomenology, naturalism and cognitive science. *Royal Institute of Philosophy Supplement*, 72, 135–167.
- Wheeler, M. (2015). The revolution will not be optimised: radical enactivism, extended functionalism and the extensive mind. *Topoi*, 1–16. doi:10.1007/s11245-015-9356-x.