CrossMark

REC: Revolution Effected by Clarification

Daniel D. Hutto[1]

Abstract This paper shows how a radical approach to enactivism provides a way of clarifying and unifying different varieties of enactivism and enactivist-friendly approaches so as to provide a genuine alternative to classical cognitivism. Section 1 reminds readers of the broad church character of the enactivism framework. Section 2 explicates how radical enactivism is best understood not as a kind of enactivism per se but as a programme for radicalizing and consolidating the many different enactivist offerings. The main work of radical enactivism is to RECtify, existing varieties of enactivism and other cognate approaches so as to strengthen and unify them into a single collective that can rival classical ways of thinking about mind and cognition. Section 3 shows how even seemingly non-enactivist explanatory offerings—such as predictive processing accounts of cognition—might be RECtified and brought within the enactivist explanatory fold. Section 4 reveals why, once RECtified, enactivist offerings, broadly conceived, qualify as genuine and revolutionary alternatives to classical ways of understanding cognition.

Keywords Enactivism · Radical enactivism · Predictive coding · Mental representations · Mental content

✉ Daniel D. Hutto
ddhutto@uow.edu.au

[1] Faculty of Law, Humanities and the Arts, School of Humanities and Social Inquiry, University of Wollongong, Wollongong, NSW 2522, Australia

Well you can twist and shout
Let it all hang out
But you won't fool the children of the revolution
—Marc Bolan, T-Rex

1 Broad Church Enactivism

Enactivism sees mind and cognition as irreducibly interactive in character. Active engagement with things and others is held to be the true basis of the psychological and epistemic situation of all cognitive beings, us included. Reversing a familiar order of explanation, the basic character of cognition is relational and dynamic, and not primarily—if at all—a matter of representing features of the world. By enactivist lights, cognition does not stop short of wide-ranging worldly engagements with what is being perceived or thought about. Cognition is not primarily a heady, brainbound affair of manipulating representations. Perceiving, a paradigm case of cognitive activity, takes time to unfold and makes direct contact with the objects with which it deals, making it a spatially and temporally extended business. From this perspective, understanding cognition cannot be reduced to knowing what occurs in brains, it also requires knowing what cognizers are doing with their bodies when adjusting to and engaging with features of the world that are afforded to them.

So construed, enactivism is a distinctive philosophical framework for thinking about minds. It is not a single well-defined empirical theory or even a particular set of such theories. It offers a broad church vision of the nature of mind—one capable of housing a diverse family of approaches. Those that adhere to the basic tenets set out above and which employ explanatory concepts such as,

Springer

autopoiesis, emergence, agency, sensorimotor contingencies, organism–environment couplings, action-perception loops, dynamical systems are all clearly in the enactivist family.

After several decades since its articulation, the basic enactivist outlook has taken firm root and is thriving in various sciences of the mind. Enactivist thinking has launched many constructive—fertile and fecund—research programmes focusing on an extremely diverse array of phenomena. This is especially true of those enactivist approaches that make use of the explanatory apparatus of dynamical systems theory. Notably, these approaches have provided "a fresh perspective on many foundational problems in cognitive science, including perception–action, memory, word recognition, decision making, learning, problem solving, and language" (Riley and Holden 2012, p. 593).[1]

In short, as a general framework, enactivism has proved productive in ways that justify taking its basic philosophical outlook quite seriously.[2] Indeed, it appears:

> There is no abating the ever-increasing popularity and influence of enactivism both in philosophy and cognitive science. With a steady flow of important publications and regular conferences it now deservingly demands serious attention as a theoretical alternative to mainstream cognitivist accounts of mind and cognition. It should then come as no surprise that with this increase in influence and popularity various distinct yet related takes on enactivism have emerged (De Jesus 2015).

The most prominent and well-developed versions of enactivism are Autopoietic-Adaptive Enactivism (AE) and Sensorimotor Enactivism (SE). Both have made important contributions to new ways of thinking about cognition.[3]

AE promotes a thoroughly biological vision of cognition grounded in a life-mind continuity thesis. It conceives of mind and cognition as emerging from the self-organizing, self-creating and self-preserving activities of a sub-set of living organisms that exhibit agency, understood in a particular way. By AE lights, cognition is inescapably bound up with the world engaging, life-preserving activity of the sort in which agents, conceived of as complex assemblies of response systems, possess a certain kind of autonomy. On the one hand, AE recognizes that the "environment plays a fundamental role on the constitution of agency, because processes in which the agent as a whole interacts with its outer environment also contribute to the maintenance of the agent" (Heras-Escribano et al. 2014, p. 3). On the other hand, AE insists that organisms are not slavishly—or mindlessly—responsive to external factors; they do not do so in a fixed or purely mechanical way.

Agents are shaped by habit but remain responsive to their current contexts. They always have some greater or lesser room for maneuver in any given situation. Consequently, even though an agent's cognitive activity is always world engaging and world relating, it is always to some degree flexible and spontaneous, and not determined or dictated in any direct way by the world. Of course, such flexibility comes at a price: even minimal freedom incurs risk. More or less effective—better or worse—couplings or engagements, relative to the needs of the agent, are always live possibilities. In acting autonomously in this minimal sense agents are always in a precarious position as they work to ensure they keep themselves "unified and distinct" (Di Paolo 2005, p. 434; see also De Jaegher and Di Paolo 2007).

Fans of AE make much of the fact that organism–environment couplings can be more or less effective. They see this as implying the existence of a kind of biological normativity that goes beyond any norms which can be associated with mere autopoiesis. Arguably, the self-organizing, self-sustaining activities of all living beings, as exemplified by the metabolic self-production of single-cell organisms, entail the existence of very basic goals tied to maintaining continuing identity. Yet it has long been recognized that any biological norms connected with such continuance are too open-ended to account for the way agents target and respond to specific features of their worlds. Purely autopoietic versions of enactivism need augmenting in order to account for the sort of biological normativity that is a hallmark of even the most basic kind of cognition (Di Paolo 2005). The kinds of biological norms associated with autopoietic activity are simply too

---

[1] Riley and Holden (2012) point out that dynamic systems approaches have been successfully employed in the investigation of these many and varied cognitive phenomena. This includes work by Kelso (1995), Turvey (1990), Beer (2009), Cadez and Heit (2011), Colangelo et al. (2004), Holden et al. (2009), Rueckl (2002), Wijnants et al. (2012a, b), Araújo et al. (2006), Busemeyer and Townsend (1993), McKinstry et al. (2008), Dale et al. (2008), Phattanasri et al. (2007), Stephen and Dixon (2009), Elman (1995), Tabor (2002). The list is indicative not exhaustive: it merely scratches the surface.

[2] This is surely so if we assume that "scientists typically know what they're talking about when they are talking about their science" (Shapiro 2014b, p. 74). Still, there are many important and philosophically interesting questions to resolve about the scope of enactivism and the kinds of explanation it offers [For a discussion of how dynamical explanations can be thought of as mechanical explanations see Zednik (2011)].

[3] These varieties of enactivist surely have had the greatest number of adherents. Proponents of AE include, e.g. Jonas (1968), Maturana and Varela (1980), Di Paolo (2005), Thompson (2007), Barandiaran et al. (2009), Barandiaran and Egbert (2013). Proponents of SE include,

Footnote 3 continued
e.g. O'Regan and Noë (2001), Noë (2004), Cooke and Myin (2011), O'Regan (2011).

weak and open-ended to capture what makes even the most primitive forms of cognition normative. Simply put, such norms can be satisfied without needing to know more precisely how this was achieved. As a result, "self-constitution of an identity can thus provide us only with the most basic kind of norm, namely that all events are good for that identity as long as they do not destroy it" (Froese and Di Paolo 2011, p. 8).[4]

Something stronger than autopoiesis is required in order to understand the normativity of cognition. Enactivists of the AE stripe think adaptivity can do that work. Adaptivity is the process of agent–environment couplings through which certain sensorimotor loops come to be favoured as "more useful for the agent than others, such that those become more salient or meaningful" (Heras-Escribano et al. 2014, p. 3). Importantly, "adaptive regulation is an *achievement of the autonomous system's internally generated activity* rather than merely something that is simply undergone by it" (Froese and Di Paolo 2011, p. 9, emphasis added).

In sum, although autopoiesis provides a necessary foundation for agency, it is the capacity of agents to adapt selectively to specific features of their environment that puts the real meat on its bones. Perhaps AE's great achievement has been to provide a way of thinking about the basic goal-directed cognition of agents in terms of biological norms without invoking any of the standard equipment that cognitivists insist is required for making sense of that phenomenon (e.g. mental contents, prior intentions, directions of fit, and so on).

SE, by comparison, has focused on understanding a smaller range of cognitive phenomena: laying stress on the integral connection between perception, action and perceptual experience (Hurley 1998; O'Regan and Noë 2001; Hurley and Noë 2003; Noë 2004, 2009, 2012). SE views perception as "a mode of activity involving practical knowledge about currently possible behaviours and associated sensory consequences. Visual experience rests on know-how, the possession of skills" (O'Regan and Noë 2001, p. 946, emphases added). In defending this line, SE rejects the idea that we form a rich and detailed inner representation when we perceive. Its enactivist commitments are clearest when it stresses the ways in which perceiving strongly supervenes on or is constituted by temporally extended, interactive worldly engagements; when it defends the idea that perception "isn't something that happens in us, it is something we do" (Noë 2004, p. 216). Accordingly, activity in neural substrates is

necessary but not sufficient for perceiving or having perceptual experience—perceiving is "realized in the active life of the skilful animal" (Noë 2004, p. 227). SE's greatest achievement has been to demonstrate the empirical robustness and explanatory power of understanding perception through its enactivist lens.

2 RECtification

What of radical enactivism then? What kind of enactivism is it? What does Radically Enactive Cognition, REC, offer that differs from other members of the enactivist family? The question is ill posed. REC is not an alternative version of enactivism with distinct explanatory tools in its own right. Technically, Shapiro (2014a) is correct to say that *on its own* REC does "little to account for the stunning successes of cognitive science" (p. 214). That is, however, to miss the point since REC never stands alone. Its analyses and arguments are designed to cleanse, purify, strengthen and unify a whole set of existing anti-representational offerings. REC's aim is *to radicalize* existing versions of enactivism and related explanatory accounts through a process of philosophical clarification.

This is most evident in REC efforts to show that SE is best and most coherently formulated in non-representationalist terms. For example, REC has strived to show that the embodied know-how of perceivers—their mastery of sensorimotor contingencies—should not be taken to imply that they or their brains possess and use a neurally-based set of rules and representations (Hutto 2005, 2011; Hutto and Myin 2013). Against a conservative reading, REC promotes a reading of SE according to which the laws of perception can be 'read off' from the activity of perceptual systems as they respond to different types of objects while denying that the work of perceptual systems requires encoding or representing such laws at any level.[5]

REC aims to show that existing enactivist approaches can account for basic cognition, without residue, by understanding it in terms of thoroughly relational, interactive, dynamically engaged, world-relating activity—activity that does not involve relating to or manipulating any kind of informational or representational content. REC seeks to clarify the true character of the explanatory resources of various enactivist and enactivist-friendly offerings in the sciences of the mind. It aims to show that making free use of the notions of information, algorithm, and representation only tends to obfuscate our accounts of cognition. It recommends that these notions should be

---

[4] Or as, Heras-Escribano et al. (2014) put it "everything that does not result in a loss of organization could be valued as something good (or at least 'not bad') for the organism" (p. 3).

[5] In this REC agrees with Burge (2010) when he observes that: "To perceive, individuals need not represent their own states or operations, even 'implicitly'" (p. 405).

either eliminated from our explanatory frameworks or otherwise radically reconceived, to a more accurate understanding of the nature and the roles they actually play in cognition.

For example, REC's uncompromising reformist zeal is motivated by the recognition that what many philosophers and scientists call 'informational content' is in fact no kind of content at all. This observation is based on two important facts: (1) that the only scientifically respectable notion of information is that of nomic covariance of some sort; and, (2) that nomic covariance isn't any kind of content. Even diehard representationalists are prepared to admit this, but they fail to draw the full consequences of these admissions.

If covariance isn't any kind of content then any science of the mind committed to explanatory naturalism that employs the notion needs to supply another scientifically reputable candidate for informational content. The Hard Problem of Content must be answered or talk of informational content cannot be taken as anything more than convenient *façon de parler*. If this problem cannot be addressed then we have no reason to believe anything backs up claims that cognitive systems are, in essence, importantly unlike merely physical interacting systems because they literally process informational contents (Clark 2008, p. 26). Moreover, failure to address the Hard Problem of Content scuppers even the most promising naturalistic theories of representational content. Close analysis reveals that failure to supply a naturalistic account of informational content undermines the foundations of such theories or forces them to radically reform (See Hutto and Myin 2013, Ch. 4).

REC sees this as a positive opportunity. It seeks to salvage some of the core ideas from teleosemantics—the most promising naturalistic theory of content to date—by putting them to a different theoretical use within the enactivist framework. The teleosemantic apparatus, REC holds, can be used to explicate a workable account of contentless basic intentional directedness as opposed to a robust semantic theory of content.

This, in turn, enables REC to provide an account of cognition in terms of active, informationally driven, world-directed engagements, where a creature's current tendencies for active engagement are shaped by its ontogenetic and phylogenetic history. Basic minds target, but do not contentfully represent, specific objects and states of affairs. Fundamentally, cognition is a matter of sensitively and selectively responding to information but it does not involve picking up and processing informational content or the formation of representational contents.[6]

This account of target-focused but contentless Ur-intentionality provides those working in the enactivist framework with a powerful tool. It is REC's major contribution: for supplying this tool enables enactivists and others to make a clean and radical break with intellectualist traditions.

This type of RECtifying clarification is needed even for enactivist approaches that most openly set their faces against representationalism. Consider that proponents of AE characterize the sort of adaptive responding required for basic cognition as a kind of 'sense making'. Sense making is said to occur when an agent treats the perturbations it "encounters during its ongoing activity from *a perspective of significance* which is not intrinsic to the perturbations themselves" (Froese and Di Paolo 2011, p. 9, emphasis added). Sense making is characterized as a "*process of meaning generation* in relation to the *concerned perspective* of the autonomous system" (Froese and Di Paolo 2011, p. 7, emphasis added). Sense making "is the enaction of a *meaningful* world by an autonomous system" (Froese and Di Paolo 2011, p. 7, emphasis added).

The so-called intrinsic meaning that is generated through sense making is neither a feature of external environment nor something internal to the agent where either is understood in isolation. This is famously illustrated by the now familiar example of a bacterium engaging in sense making and thus enacting its world by responding in different orientations to a sugar gradient. Having the status of an affordance, the "sugar's edibility is not an intrinsic property: it is only valuable in relation to the agent that takes advantage of it" (Heras-Escribano et al. 2014, p. 4; Thompson 2007, p. 125). This meaningful property only comes into being when the agent relates to features of its world. Still, even though the meaning in question is thoroughly relational, crucially according to AE, the relata cannot be individuated or characterized independently of the agent's engagements with *its* world (Thompson 2007, p. 74).[7]

It is problematic to assume that basic minds are capable of 'sense making' and 'meaning generation' in anything like the robust sense implied by the standard connotations of these terms. Although there is no reason to deny cognitive status to non-contentful world-directed activities of

---

et al. (2014), it does not seek to develop "new ways of understanding informational content" (p. 2).

[7] Putting all of this together, defenders of AE take it to be "appropriate to consider adaptive autonomy as the most basic form of life, and sense-making as the most basic process of living" (Froese and Di Paolo 2011, p. 9). Thus: "If autopoiesis (or autonomy) suffices for generating a 'natural purpose' (Kant 1790), adaptivity reflects the organism's capability — necessary for sense making — of evaluating the needs and expanding the means towards that purpose" (Froese and Di Paolo 2011, p. 9).

---

[6] REC shows that enactivists can get by without having to fall back on a notion of informational content. Thus, *pace* Heras-Escribano

living creatures—including ourselves—being capable of detecting, tracking, and interacting with salient features of an environment alone does not suffice for 'sense' or 'meaning' making understood in any standard sense. Having a mind that 'makes sense' of its world and generates meaning about it is not a fundamental biological endowment. That ability requires special forms of scaffolded engagement. In particular, it requires participating in and mastering normative practices that are beyond the reach of simple organisms.

Thus a standard verdict is that AE, in assuming that the properties of complex cognitive systems will be found even in very simple cognitive systems, "gets things the wrong way round" (Menary 2015, p. 3). The root problem is that the norms for meaning making are not of the biological sort AE identifies. Thus AE "takes normativity to permeate all biological and cognitive functions. But … this supposed normative character of natural reactions fails to satisfy some specific requirements for something to be classified as normative [in the robust sense required]" (Heras-Escribano et al. 2014, p. 8). More precisely, the complaint is that the kind of biologically based norms exhibited by active "regulations with the environment, instead of being conceptual, are normative but not contentful" (Heras-Escribano et al. 2014, p. 10). For this reason the sort of misalignments and failures of world engagement that can occur at the level of basic cognition do not involve making errors that are anything like errors of contentful judgement. Failures to engage with the world effectively are not, and are not explained by, failures to describe, depict or say how things stand with the world.

REC's biosemiotic account of contentless intentionality is useful at just this juncture. It provides a way to make sense of the kind of biological normativity associated with basic cognitive activity while at the same time allowing that "the best explanation of the origins of norms [of the sort needed to speak of content] is that they are natural phenomena located at the social level" (Heras-Escribano et al. 2014, p. 10; see Hutto and Satne 2015).

This is a shining example in which REC's alternative account of contentless informational sensitivity and targeted intentionality provide a solid means for enactivists to put their positive explanatory offerings on a stable theoretical footing. Rather than offering new and different explanatory tools in addition to those provided by other forms of enactivism, REC's aim has always been to conduct the philosophical work needed for developing and refining such tools to provide a set of genuine alternatives to standard explanatory equipment employed by classical, intellectualist cognitive science.

RECtification is needed if we are to see a true "crystallization of enactivism" and the full and complete development of "a positive alternative to representationalism"

(Heras-Escribano et al. 2014, p. 1). Nor should RECers shy away from incorporating other enactivist-friendly approaches within the general enactivist framework. For example, Gibson's groundbreaking work on affordances—which has long provided a positive explanatory basis for a non-representationlist cognitive science, prior to enactivism's maturation—is also ripe for RECtification.[8]

Ecological psychology makes use of the notions of direct perception of environmental affordances (Rietveld 2008; Rietveld and Kiverstein 2014; Kiverstein and Rietveld 2015; Chemero 2009). Classical ecological psychology regards perception as an active, relational phenomenon; organisms as sets "of abilities"; and niches "as the set of situations in which one or more of [an organism's] abilities can be exercised" (Chemero 2009, pp. 147–148). Chemero (2009) seeks to update ecological psychology, making it "dynamical root and branch" (p. 150). His aim is to provide more refined theoretical tools that can best serve the experimental and explanatory needs of ecological psychologists. This requires giving pride of place to organismic interactions and how they develop cognitive tendencies over time.

Ecological approaches, especially those promoted under the radically embodied cognition banner, are natural allies for enactivist approaches. Indeed, it has been suggested that enactivism and ecological psychology have much to gain by being brought together within a larger framework. Chemero (2009) noted the theoretical value in such unification and suggested some initial steps toward it, although he rightly recognized that "much more work is required to genuinely integrate ecological and enactive approaches" (p. 154).

The REC analysis is that a thoroughgoing integration of ecological dynamical approaches requires, *inter alia*, clarifying common talk of the 'provision', 'use', 'gathering', and 'pick up' of information 'about' affordances that is prevalent in the work of some of its leading proponents (see Chemero 2009, pp. 154–161). The problem with such talk is that it suggests an underlying commitment to an information-processing story that is inconsistent with non-representationalist accounts of mind and cognition.[9]

---

[8] Other E-approaches to mind and cognition—such as Menary's (2007, 2015) Cognitive Integration Theory and Malafouris's (2013) Material Engagement Theory—seem even more amenable to RECtification.

[9] Millikan (2005) who also endorses some central Gibsonian ideas about the active nature of perception falls foul of talk of 'collecting', 'picking up', 'applying' and 'transmitting' natural information. She writes, "thinking of a substance involves the ability to recognize it, as it were, *in the flesh*, not merely the ability passively to contemplate its properties. We have thoughts of substances in order to be able to collect information about substances, which information we pick up on some occasions and apply them on others. To pick up information about a substance you must be in a position to interact with the substance, other things that are influenced by the substance or that

RECtification is therefore needed in order to clarify the notion of information and the role it plays in ecological accounts if they are to be part of a larger non-representationalist framework.

Van Dijk et al. (2015) acknowledge this tension in the very heart of Chemero's work and also in the writings of other prominent theorists working within the ecological tradition. They admit that it is "hard to get a contentless reading of even the most progressive ecological theories" (Van Dijk et al. 2015, p. 212).[10] Even so, they don't see this as an intractable problem for ecological accounts because, on their assessment, talk of organisms responding to 'information about' affordances and of 'informational pick up' does not reflect a genuine theoretical commitment of such approaches. Instead they take such talk to be nothing more than a hangover of an unreflective use of language, thus it is easily revisable. If they are right then ecological accounts can be rendered REC-friendly painlessly.[11]

Once again, what should be abundantly clear at this point is that REC does not aim to provide distinct explanatory tools of its own in addition to those supplied by other varieties of enactivism. For this reason REC does not see the various members of the enactivist family as competing rivals. Rather it aims to sanitize what such enactivist approaches have to offer, removing any residual vestiges of representationalism: REC seeks to radicalize them. That is its programme. Its ultimate aim is to clarify and unite the various non-representational approaches to cognition, demonstrating how they work together under one philosophical roof to cooperatively provide genuine, complementary alternatives to classical cognitivism.

3 Peace Through Clarification

How far can the RECtifying programme go? It might seem that some of cognitive science's explanatory tools are simply off limits to REC. Yet since appearances can deceive, it would be a mistake to rush to judgment: careful investigation is needed on a case-by-case basis.

Consider what might appear to be a hard case, RECtifying the theory of cognition that trades under the names Predictive Coding, Prediction Error Minimization and Predictive Processing (PPC for short). PPC is causing a real stir in philosophy and neuroscience (Clark 2013a, b; Hohwy 2013, 2014; Friston 2010; Friston and Stephan 2007). PPC's leading idea is that the true, indeed *only*, work of brains—their ceaseless cascade of multi-level, multi-layered cortical processing—is all part of a singular effort to predict sensory deliverances.[12] PPC represents a dramatic reversal of traditional cognitivist thinking. It regards the core business of cognition to be making pro-active, probabilistic, Bayesian predictions about likely sensory perturbations as opposed to constructing internal models of the world that are built upon passively received information furnished by the senses.

Prediction error occurs when there is a mismatch between what brains predict and what is supplied to them by the senses. The brain's aim is to minimize the divergences between what it anticipates and what is sensed. This can be achieved either by making better guesses or making adjustments so as to get more fitting sensory inputs. This is known as the reduction of prediction error or uncertainty.

Some deem PPC as heralding a quite profound, radical sea change in our thinking about the mind. For example, Clark (2015a) claims, "Predictive processing plausibly represents *the last and most radical step* in [the] retreat from the passive input-dominated view of neural processing" (p. 2, emphasis added).[13] Dramatically, he observes that if PPC is along the right lines then "just about every detail of the passive forward-flowing model [as promoted by classical cognitivism] is false" (Clark 2015a, p. 2). Instead, if PPC is right, cognition has a fundamentally

---

Footnote 9 continued

influence it. Natural information is transmitted in the causal order, and you have to be in the causal order, with whatever the information is information about to receive it" (Millikan 2005, p. 115).

[10] Van Dijk et al. (2015) propose that a more thoroughgoing analysis of Gibson's views reveals that, even though "the notion of 'information pick-up' … takes on a content-carrying connotation from Gibson's early work, [but it can be] understood in a content-less sense. Having the sensitivity, or the openness, to 'resonate' to the ambient patterns available, the animal picks up [on] those patterns as information for perceiving and acting. Such a reading, we feel, would give a fruitful and more charitable account of ecological theories" (p. 213). Crucially, in RECish spirit, they note "there need not be any content involved at all, as information for affordances cannot be evaluated as being more or less true or accurately corresponding to an affordance – there are no conditions to satisfy it being about the affordances… information can be more or less useful for adapting to the environment, that is all" (p. 213).

[11] Following REC's lead, Van Dijk et al. (2015) realize that ecological psychologists need to embrace the notion of 'information for' in favour of the notions of 'information about' (p. 212). In contrast these authors propose that "Information needs to be understood 'teleosemiotically' … the high level of array-environment correspondence makes the patterning in the array very useful to the person. But only as these patterns are used, need they be considered information for perceiving or acting on affordances" (p. 213).

[12] Hohwy (2014) reports that, according to PPC, "prediction error minimization is the *only* principle for the activity of the brain" (p. 2, emphasis added).

[13] Clark (2015a) identifies a number of 'quite radical' implications of PPC: (1) the core flow of information is top-down—the forward flow of information is replaced by the forward flow of prediction error; (2) motor control is just top-down sensory prediction; (3) efference copies are replaced by top-down predictions; (4) cost functions are absorbed into predictions (p. 3).

"'restless', pro-active, hyperactive and loopy character" (Clark 2015a, pp. 1–2).

This last observation might make it seem as if PPC is naturally suited for accommodation within an enactivist framework.[14] It raises the question: might PPC be a suitable target for RECtification after all? Exploring this possibility in full detail requires more careful analysis than can be provided in the limited space available here. Even so, it is worth making some initial, first pass observations to establish not only that RECtification isn't out of the question but also why it ought to be welcomed by proponents of PPC. Making this case, even in sketch, both prepares the ground for pursuing the question further and has the added advantage of providing another vivid example of what RECtification looks like in action.

3.1 Secluded Minds Versus Open Minds

The prospect of RECtifying PPC is a tantalizing challenge precisely because several recent papers regard PPC as absolutely wedded to the idea that the brain trades in contentful representations. For example, Hohwy (2014) assumes that the brain's predictions about likely sensory input "necessarily rely on internal representations of hidden causes in the world (including the body itself)" (p. 17).

Why assume this must be the epistemic predicament of brains? Apparently, it is because it would be an "ideal but impossible design" for the brain to make any direct comparison between its internal estimates and "true states of affairs in the world" (Hohwy 2014, p. 4). As Hohwy rightly observes the brain itself is in no position to compare what it represents with what is so represented. Given this access problem, the best a brain can do is make inferences to the best explanation about how things stand with the world. In the best case, the brain hits on a hypothesis that best explains away the occurrence of some evidence and is, thus, self-evidencing.

Importantly, Hohwy (2014) notes that "the notion of self-evidencing appears to be the epistemic cousin to the dynamic systems theory notions of self-organization and self-enabling, which are often used to explain enactivism" (p. 19). The major difference is that in assuming representationalism from the start, Hohwy's (2014) take on PPC paints "a picture of the brain as a secluded inference-machine" (p. 19). It is because Hohwy assumes representationalism that he advocates "decoupling the brain from the body and the environment in an epistemic sense" (pp. 18–19). The mind-brain is forever secluded and cut-off from knowledge of the world. This is to assume a "stark

mind-world schism" that entails global scepticism (p. 19). In pressing this line, Hohwy (2014) also rejects a halfway house approach of the kind Clark offers in which notions of coupling are used in order "to argue in favour of a less secluded, more open mind-world relation" (p. 20).

How plausible is the self-evidencing view of PPC? Doubtless sharp-eyed philosophers will observe that having a global skeptic forever on one's back will be hard-going for any creature in the 'getting it right' in order to 'act successfully in the world' business. Even sharper-eyed philosophers will wonder how epistemic states of mind that are in principle secluded from the world could ever come to have contents that refer to, or are about, inaccessible hidden causes that they putatively represent in the first place. Hohwy's self-evidencing proposal appears to run straight into well rehearsed problems about how, say, a brain-in-a-vat could possibly form the thought that it is a brain-in-a-vat, or for that matter, manage to harbor thoughts about any external topics.[15]

Why on Hohwy's model isn't the brain restricted to thoughts about sense data? There is a great deal of epistemic security for the brain if it only makes predictions about an accessible, sensory world. It can't be wrong about the state of the external world if it can't even think about such a world. But this sort of epistemic gain comes at the high price of limiting the brain to thinking only about non-worldly topics.[16]

Casting PPC in such a restricted epistemic light puts it in the same boat with those extreme forms of idealism that call the very idea of an external world into question; those which are at odds with metaphysical realism. As such, it is difficult to understand how Hohwy's policy of total epistemic seclusion squares with his assumption that the brain trades in contentful representations. *Prima facie*, the idea that brains are hermetically sealed in an epistemic sense conflicts with the idea that the content of neural representations can be understood in terms of reference to external items and truth conditional judgments about such items of the sort associated with a standard correspondence theory of truth.

In sum, ditching a secluded brain reading of PPC has the serious advantage of obviating the need to address age-old epistemological and pyschosemantic problems (at least at the level of basic perception where we make first contact

---

[14] Indeed Clark (2015b) thinks PPC should be embraced by enactivists because it bears special gifts—namely, it supplies the explanatory resources to "cash … enactivist cheques" (p. 3).

[15] As Putnam pointed out, long ago, "Although the people in that possible world [the brain-in-a-vat world] can think and say, they cannot … refer to what we can refer to. In particular, they cannot think or say that they are brains in a vat (even by thinking 'we are brains in a vat')" (Putnam 1988, p. 8).

[16] The secluded brain hypothesis suffers from the same problem as old fashion theories of AI. They need to supply "a theory of language-and-the-world, whereas, in fact, [they provide] only a theory of language-and-the-insides-of-the-machine" (Fodor 1981, p. 209).

with the world). Developing a credible alternative reading would require giving a full reply to Hohwy's (2014) 'basically friendly challenge' of showing how enactivism "can avoid an epistemic, inferential reading in terms of the self-evidencing that entails an evidentiary boundary and thus decoupling" (p. 19). Answering that challenge is a job for philosophy. Ultimately, RECtification is required. We will come to that.

3.2 Compromise and Collaboration?

Clark (2015b) has argued that PPC and enactivism can mutually illuminate each other in a way that could mark the end of the representation wars. His paper predicts, though does not guarantee, peace in our time. As Clark sees it, situating PPC properly ought to provide a long awaited olive branch for brokering peace between representationalists and enactivists. Why so? Peace is allegedly on the cards because although PPC "openly trades in talk of inner models and representations, [it only] involves representations that are action-oriented through and through" (Clark 2015b, p. 4). Consequently, the representations of PPC "*aim to engage the world* rather than depict it in some action neutral fashion" (Clark 2015b, p. 4, emphasis added).

Hence, even though PPC is heavily committed to internal models, "instead of *simply describing* 'how the world is', these models—even when considered at those 'higher' more abstract levels—are geared to engaging those aspects of the world that matter to us. They are delivering a *grip* on the *patterns that matter* for the *interactions that matter*" (Clark 2015b, p. 5, emphases added). So, as Clark presents the situation, "What is on offer is thus just about maximally distant from a passive ('mirror of nature') story about the possible fit between model and world" (2015b, p. 4).

Clark's plan for peace clearly requires a bit of give and take from both sides—a bit of compromise. Indeed, Clark cannot see another way forward. For despite recognizing the power of enactivist explanations that he sees lying at the heart of the PPC story, he has difficulty seeing how to tell that story "in entirely non-representational terms" (Clark 2015b, p. 5).

It is easy to see what motivates this assessment. Clark offers an analogy to provide an intuitive sense of the brain's situation when making active inferences. He asks us to imagine a game in which one participant attempts to describe what a second participant is seeing while the latter moves through a familiar environment—the living room of the first player's house. The catch is that the first player has no direct access to the visual scene and so can only make best guesses about what the second player is likely to see.

The second player's role is to speak up and correct those guesses should they go awry and to remain silent otherwise. Hence if player one says "There's a vase of yellow flowers on the table in front of you", the second player will either deny this or remain quiet.

Moving beyond analogy, how should we understand this at the level of theory? When it comes to understanding sense deliverances (represented by the second player's contribution), Clark (2015a) tells us that:

> *in a very real sense*, the prediction error signal is not a mere proxy for incoming sensory information – it is sensory information … your 'error signal' carried some quite specific information … *the content* might be glossed as 'there is indeed a vase of flowers on the table in front of me but they are not yellow'. This is a pretty rich message. Indeed, it does not (content-wise) seem different in kind to the downward-flowing predictions themselves. Prediction error signals are thus richly informative (p. 5, emphasis added).

The trouble is that anyone hoping to explain what the senses deliver in this way faces a hard choice: Either take talk of rich informational messages and content seriously (and pay for it by answering the Hard Problem of Content) or go radical, ditch the idea that the information in question is contentful, and significantly revise this story.

It might be thought that there is an obvious, better way to go for those who think representations are action-focused through-and-through other than try to solve the Hard Problem of Content directly (by showing how mere covariation adds up to content). The natural move at this juncture would be to make appeal to Action Oriented Representations, explicating these through the lens of Millikan's teleosemantic theory of content. Action Oriented Representations can be understood as Pushmi-Pullyu representations (for an updated rehearsal of the rationale behind this move see Clowes and Mendonça 2015, pp. 3–5).

Notably, Pushmi-Pullyu representations, although primitive, possess not one but two kinds of content—both descriptive and directive content (see Millikan 2005, pp. 173–175). This being so to think of Action Oriented Representations in such terms is to think of them as only maximally unlike the passive 'mirror of nature' representations in that they are not passive. In other respects they are as representational as representations can get. Indeed, this is why for Millikan (1993) Pushmi-Pullyu's provide the

---

[17] Of course, none of this should come as any surprise. Millikan's aim was always to revive Wittgenstein's so-called Tractarian picture theory (see Millikan 2005, Ch. 4). Her aim was to fill in the gaps of Sellars' project. Thus she was always pursuing a quite different philosophical agenda than that of answering the explanatory needs of cognitive science. See Hutto (2014) for a discussion.

primitive ground for "flatfooted correspondence views of representation and truth" (p. 12).[17]

Understanding the content of such representations by appeal to proper functions is to focus on the effects a device is *supposed to* or *meant to* produce. History sets the standard. It is history that determines whether a current state represents correctly or incorrectly. As a consequence, content, understood via biosemantic theory is debarred from playing any kind of causal role in the synchronic production of intelligent behavior.

Representational status is conferred on states by being based not on "what they do but *why* they work" (Millikan 2005, p. 97). Descriptive representations work by bearing a correspondence to what they represent. On this scheme, representational content is patrician not plebian in character. In other words, content never gets it hands dirty. Mental content understood in terms of proper function does no mechanistic work. For teleosemanticists, the focus of attention is on ultimate as opposed to proximate explanations, on structuring as opposed to triggering causes (Dretske 1988, chs. 1 & 2). This is why Millikan's normal explanations are concerned to specify the historical conditions by which cognitive devices were selected; such explanations do not seek to answer questions about how such devices operate in the here and now.

Millikan (1993) makes this abundantly clear: "having a certain history is not, of course, an attribute that has 'causal powers' … that a thing has a teleofunction is a causally impotent fact about it" (Millikan 1993, p. 186). To adopt a teleosemantic account of mental content is to forego the idea that mental contents can possibly feature in the mechanistic or causal explanations. When we make appeals to content we must focus on what a device is supposed to do, not what it is disposed to do.

The upshot is that appealing to teleosemantics is at best a means of explaining how Action Oriented Representations get their content but it is in direct conflict with the demand that representational explanations in cognitive science are supposed to "answer how-questions about cognitive capacities, and not … why-questions about particular behaviors or actions" (Gładziejewski 2015a, p. 66).[18] This being the case representational content, understood in teleosemantic terms, cannot do one of most basic jobs earmarked for it. In describing that job Shea (2013) notes:

What adverting to content does achieve, however, is to *show how the system connects with its environment:* with the real-world objects and properties with which it is interacting, and with the problem space in which it is embedded. The non-semantic description of the system's internal organisation is true of the system irrespective of its external environment. Content ascriptions help explain *how it interacts with that environment* (p. 498, emphases added).

A teleosemantic theory of content cannot answer Shea's question because at most it can only establish that a system connects to and targets certain features of the environment; it is in no position to explain how systems do so. And if it can't answer Shea's question then it cannot solve the problem Clark faces—it cannot explain how content might be literally supplied to the brain via the senses. Worse still, it is open for us to wonder if the connections of which Shea (2013) speaks are best characterized in semantic terms anyway, at all. Why should having a bio-history confer robust semantic status on internal states, even those that determinately target particular worldly offerings?[19]

RECers have long argued against this supposition. Nor are they alone in thinking that "Millikan has not provided an adequate theory of content. Millikan's technical apparatus does define a relation that can hold between a system's mental sate and properties sometimes instantiated in the environment. But … the relation so defined is not 'has as its content that'" (Pietroski 1992, p. 268). Others too think that biological functions can explicate "an important kind of natural involvement relation … [but] not … representation or anything close to it" (Godfrey-Smith 2006, p. 60). The intentional directedness of basic cognition need not be cashed out in semantic terms but might be understood, much more weakly, as instantiating "some kind of privileged relation" (Rupert 2011, p. 101).

The punch line is that positing Action Oriented Representations understood via the lens of teleosemantics rules out the very idea that mental contents might causally drive actions. Nor is it obvious there is any ground for thinking that these so-called 'representations' possess any kind of content at all. Bringing this back to PPC, for all of these reasons there a number of serious problems with Clark's answer to this question:

What are the contents of the many states governed by resulting structured, multi-level, action-oriented, probabilistic generative models? It is … precision-weight estimates … *that drive action* … such looping complexities … make it even harder (perhaps impossible) adequately to capture the contents or the

---

[18] On the standard view: "Information processing theories effectively offer a wiring diagram showing how inputs affect states of the system and, in conjunction with other states of the system, issue in behavioural outputs. What does it add to that wiring diagram to label various nodes with representational contents? A realist about mental representation is committed to the reality of the internal particulars described in the theory, and of their contents" (Shea 2013, p. 498).

[19] For an extended argument why we should abandon this idea see Hutto and Myin (2013), Ch. 4.

cognitive roles of many key inner states and pro-
cesses using the vocabulary of ordinary daily speech
(Clark 2015b, p. 5, emphasis added).

3.3 Maps and Models

Let us consider one last way of possibly keeping repre-
sentationalism in play. Despite recognizing the limits of
teleosemantic approaches some authors, such as
Gładziejewski (2015b), insist that the representational
pretensions of PPC are entirely justified. Indeed, it has been
argued that in the final reckoning PPC might be as repre-
sentational as cognitive-scientific theories get.

Why think so? Gładziejewski (2015b) is very clear that
he assumes "prediction error minimization aims to mini-
mize the mismatch between how things are and how the
brain/mind 'represents' them as being".[20] If this is right
then PPC must operate with the strong notion of repre-
sentational content—the generic notion that is used at large
in the classical cognitive science literature (for an excep-
tionally clear discussion of what this entails see Travis
2004, esp. pp. 58–59). To understand representational
content in this way is to subscribe to the idea that repre-
senting the world contentfully is a matter of taking ('rep-
resenting', 'saying', 'asserting', etc.) that the world is a
certain way such that it may or may not be that way.
Representational content therefore implies correctness
conditions of some kind, which can be variously construed
as truth, accuracy or veridicality conditions.

The assumption is that the more accurate the brain's
generative model is in terms of its "likelihoods, dynamics,
and priors" the more accurate its hypotheses will be about
the causal–probabilistic structure of the external world. The
assumption motivates thinking that representational con-
tents of a quite traditional kind must play a central part in
PPC explanations. To illustrate the point, Gładziejewski
(2015b) imagines a case in which a human brain uses a
less-than-accurate generative model of the world and thus
"settles on the hypothesis that it is seeing a plush imitation
of a tiger … when what it in fact faces is a live tiger".
Getting it wrong in this sort of case has fairly obvious
costs. And even if the brain has no direct access to tigers
via perception we can expect that when "what one is in fact

observing is a tiger, then the hypothesis that the inflow of
sensory information has been caused by a tiger will gen-
erate (on average) a smaller prediction error than alterna-
tive hypotheses—including any that attribute the causal
origins of the incoming signal to a plush toy or a domestic
cat" (Gładziejewski 2015b).

It is clear to see why anyone who understands the role of
perception in this way will think that PPC necessarily
trades in explanations involving representational content.
The real question is, in the light of previous failed attempts,
how does Gładziejewski (2015) intend to pay for the rep-
resentational content that PPC explanations putatively call
upon?

Gładziejewski (2015b) proposes that PCC should pos-
tulate internal representations whose functional profile is
nontrivially similar to the functional profile of cartographic
maps. He identifies four features that qualify such maps as
representations: allegedly they represent by: (1) struc-
turally resembling features of some domain; (2) guiding the
actions of their users; (3) doing so in detachable ways (e.g.
they can be used 'off-line'); and (4) allowing their users to
detect representational errors.[21] Drawing on work by
O'Brien and Opie (2004), Gładziejewski (2015b) presents
the basic idea in this way maps "represent in virtue of
sharing, to at least some degree, a relational structure with
whatever they represent".[22]

Explicating this idea, O'Brien and Opie (2015) conceive
mental representation in *structural* or *analog* terms—in
terms of physical analogies holding between representa-
tional content and what they represent. Representational
contents just are intrinsic structural properties of repre-
sentational vehicles. Mental representations are structures
that share resemblance properties of some kind with what
they represent. The content of an analog representing
vehicle is fixed solely by structural resemblances holding
between its vehicle and its object. The resemblances in

---

[20] Rehearsing a view found in countless cognitive science textbooks,
Rey (2015) reports, "'representation' has come to be used in
contemporary philosophy and cognitive science as an umbrella term
to include not only pictures and maps, but words, clauses, sentences,
ideas, concepts, indeed, virtually anything that is a vehicle for
intentionality (i.e. anything that stands for, 'means', 'refers to', or 'is
about something')" (p. 171). Of course, a consequence of adopting
this 'broad usage' is that representationalists face the difficult
problem of "determining by virtue of what something represents
whatever it represents – namely has the representational content that it
has" (Rey 2015, p. 171).

[21] In making this proposal Gładziejewski (2015a, b) hopes to address
Ramsey's (2007) job description challenge, and deal with the worry
that representational constructs used in the cognitive sciences are
often representational in name alone—viz. that representational
terminology too often serves as "an empty and misleading orna-
ment". As he puts its, "It is easy to say that representations are
component parts of mechanisms that play the functional role of a
representation. But it is much harder to answer the question of what it
means to function as a representation within a mechanism. When are
we justified in attributing the role of a representation to a component
of a neural or computational mechanism? What exactly does a
component have to do within a mechanism in order to be justifiably
categorized as a representation?" (Gładziejewski 2015a, p. 67).

[22] Gladziejewski (2015a) identifies Grush's (1997, 2004) emulator
theory of representations—which holds motor control perception and
imagery make use of internal emulators of the body and world—as a
shining example of the sort of theory that posits constructs which
comfortably satisfies the four criteria required, by his lights, to count
as representations (Gladziejewski 2015a, p. 85).

question are thought to depend only on intrinsic properties of the vehicle and objects.

Going with this sort of structural-resemblance theory of representational content will be attractive to anyone who is convinced that other naturalized theories of content face intractable problems. However, notably going the analog way seems to involve substantially weakening the very idea of representational content itself, indeed to vanishing point. That's the rub. For the notion is so reduced that it becomes difficult to see in virtue of why internal structures need to be thought of as bearing contents at all (or indeed, assuming they do have contents how the contentful properties rather than other properties of the structures that allegedly bear them play any explanatory role in cognition). Put simply, once we have gone this far it is hard to see why even cheaper contentless alternative accounts of resembling structures aren't better placed to do all of the relevant explanatory labour.

To illustrate the point, assume that for some structure to function as a map certain resemblances must hold between the map and the mapped domain. That can be granted without assuming that structural resemblances entail the existence of any kind of representational content. Maps can be used to navigate an environment because certain correspondences hold. But when they are so used it does not follow that the exploitation of map-like correspondences for the purposes of navigation entails using the map representationally such that the successes and failures of such effort need to be understood in terms of representational content.

Here thinking about keys and locks is revealing. Let us assume that every mental structure has its own unique geometry, unique structural properties. Accordingly, the 'shape' of such structures and how they interact in the machinery of the mind are what drive cognition. It is easy to see how the analogue properties of a given structure plausibly "determine the causes and effects of its tokenings in much the way the geometry of a key determines which locks it will open" (Fodor 1991, p. 41). It is not obvious how any putative content a structure might bear could do likewise or indeed why such structures should be thought of as intrinsically contentful.

Assume that the structural-resemblance notion of 'representation' may have explanatory power and applicability. The money question is why, even assuming structural resemblances do important cognitive work, think that they do it in virtue of being contentful (in the sense of having correctness conditions such as truth or accuracy conditions)?

In the absence of further details it seems—borrowing from Myin and Hutto (2015)—that all the explanatory work can be outsourced to less costly employees, structural

resemblances, being actively exploited in systematic ways in order to bring actions to bear on teleologically fixed targets. The existence of the structural resemblance seems all that is needed to do all the work in explaining a system's adaptiveness—there is no need for content to play any part in this story, and no room for it to do so.

From its inception, REC has exposed the explanatory hollowness of inflated hyper-intellectualism about cognitive processes, such as perception (see Hutto 2005; Hutto and Myin 2013, ch. 5). Reprising that reasoning and focusing on vision, Orlandi (2014) shows there is a less expensive way to go in understanding perceptual processes than buying into representationalist renderings of PPC accounts. She reveals that even if we can describe perceptual processes in Bayesian terms there is no need and no advantage in characterizing the brain as literally making Bayesian inferences in carrying out its work. Orlandi (2014) provides a perfect statement of the REC take on this issue[23]:

We can explain the central phenomena that we need to explain by thinking of the visual process as mediated by functional states and features that are better understood non-representationally, and making reference to environmental conditions, in particular to statistical regularities in the world with an eye to organismic needs … a system may have features—for example, wires or constraints—that developed, and continue to develop, under evolutionary and environmental pressure. These features have a certain function. They make the system act lawfully; that is, they make the system act in a way that is describable by principles. The principles, however, are in no sense represented by the system and encoded within it (p. 3).

What moral can we draw from this? Perceiving can be understood as a dynamic, active process based on having a certain history of interactions and information sensitivities in the current context through which we can make contact with the world. Even with this use of maps and models employing structural resemblance it does not follow that such maps and models intrinsically represent the world contentfully. Maps and models may help us to engage with an aspect of the world without our forming hypotheses or making true or false, accurate or inaccurate 'claims' about how things stand with the world. What hold true for us in this case, holds equally true for our brains.

3.4 A Lasting Peace

We are now, finally, in a position to give a stronger reply to Hohwy's challenge (see Sect. 3.1). The epistemic reading

---

[23] Despite agreeing with REC about the nature of perceptual processes, Orlandi (2014) disagrees with REC in supposing that the product of such processes are plausibly representational (p. 5).

Hohwy ([2014](#)) advances is committed to thinking that perceiving involves representational content. But why assume this? Why think that perceiving entails the existence of perceptual contents with correctness conditions (truth, accuracy, veridicality conditions)? To assume perceiving must involve representational content is bound up with the idea that any instance of *bona fide* perceiving must involve taking or depicting the world to be a certain way such that it might not be that way. The assumption that Cognition must Involve Content (CIC) is the very assumption that REC denies.

To let go of CIC is to let go of the very core of the orthodox representationalist vision of cognition in a way that spares us the problem of being epistemically cut off from the world. If perceiving isn't fundamentally a matter of representing the world then there is simply no question of our perceptions getting things right or wrong in basic cases. How then to make sense of PPC's error minimization? Perhaps reducing uncertainty and reducing free energy can be cashed out in non-intellectualist terms—thus not as a matter of epistemic error—but without loss of explanatory power.

Recall that according to PPC, the brain is constantly seeking to minimize the degree of mismatch between internally generated sensory predictions and incoming sensory signals from the external environment. PPC accounts assume that there are active anticipations grounded in structural and functional changes wrought in the brain through an organism's history of interactions (Byrge et al. [2014](#); Bruineberg and Rietveld [2014](#)). Even Hohwy ([2014](#)) admits that perception can "be described as empirical Bayesian inference, *where priors are shaped through experience, development and evolution*, and harnessed in the parameters of hierarchical statistical models of the causes of the sensory input" (p. 4, emphasis added).

REC too assumes that prior anticipations and expectations are grounded in structural and functional neural changes wrought through an organism's history of interactions. But it takes anticipations and expectations to be contentless. What we do and how we do it—what we experience—leads to changes in our neural set up and what we expect to experience. Yet having expectations about what we will experience sensorily need not be thought of in epistemic terms; it involves nothing like making contentful claims about the state of the world. Nor need we think of sensory perturbations that surprise as contentful messages that contradict the content of prior expectations. Although the senses are sensitive to information in the environment, they can do their action guiding work in a strictly silent manner (Travis [2004](#)). There are ways of making sense of the function of the senses in which representational contents play no part (Akins [1996](#)). Yet even on such non-representationalist construals it is still possible to talk about what an organism expects to experience on some occasion as being in tension with—or failing to 'match'—features of its current sensory experience. This being so, our expectations can fail to match incoming sensory experience without this activity being construed as an evidence-based operation.

This conclusion follows naturally if the senses do not have the job of telling us "how things stand objectively with the world" but rather of trying to ensure—within their sub-optimal limits—that organismic activity satisfies specific, narcissistic organismic needs. Satisfying such needs surely involves being sensitive and adjusting to the "causal–probabilistic structure of the world" but such adjustments need not be evidence based and representationally driven.[24] Consequently, on this analysis we can shelve the philosophically confounding talk of the brain making contentfully based 'predictions', 'inferences' and 'hypotheses'.

To return to Clark's claim that the models in PPC do more than just describe the world, we can ask why think of models as doing any describing at all? "Why not simply ditch the talk of … internal representations and stay on the true path of enactivist virtue?" (Clark [2015b](#), p. 4). Why not, indeed! Finding our way to this straight path is not a matter of brokering a theoretical peace through compromise, it is a matter of achieving philosophical peace through clarification. After all, "The real discovery … is the one that gives philosophy peace, so that it is no longer tormented by questions which bring itself into question" (Wittgenstein [1953](#), §133).

---

[24] Going the REC way requires abandoning the all-encompassing vision of standard formulations of PPC which holds that the sole explanation of adaptive error minimization is the reduction of free energy. While the reduction of free energy principle is central to PPC, it should not be regarded as 'the' foundational, ultimate explanation of all adaptive behavior; the one factor that drives it. PPC must surrender this pretension given that the optimal strategy for reducing surprise and minimizing predictive error would be, as the 'dark room' objection highlights, to find a stable environment and engage the world as possible. Clearly, the 'dark room' strategy cannot explain why the world teems with so many diverse forms of adaptive life that employ an incredible variety of adventurous cognitive strategies. Highlighting the explanatory limitations of relying on a single principle to explain all of this, Menary (forthcoming) supplies a compelling argument for relinquishing the 'isolated brain' interpretation of PPC in favour of situating the PPC enterprise within a broader, more 'open minded' and pluralist explanatory framework—one which assumes that to explain adaptive life and cognition demands appeal to a wider set of grounding evolutionary principles and not just the idea that organisms seek to minimize free energy. This argument against the secluded brain formulation of PPC only concerns PPC's official explanatory ambitions. It is, thus, independent of the epistemic and semantic concerns raised above. Nevertheless, they fit together as a seamless package.

4 Revolution in Mind

Suppose we go the REC way. How radical is the required rethink? Enactivism is often heralded as offering a new paradigm for thinking about cognition (Stewart et al. 2010). On the face of it, this verdict seems justified if it is true that enactivism is 'gradually supplanting' its traditional cognitivist-computational competitors (Cappuccio and Froese 2014, p. 3). The verdict appears justified if it is true that enactivism has "matured and become a viable alternative" to such approaches, yielding "methodological advances" that "avoid or successfully address many of the fundamental problems" faced by their rivals (Froese and Ziemke 2009, p. 466).

In advocating the PPC-enactivist alliance, Clark (2015a) too speaks of our conception of mind being turn upside down, of 'radical' conceptual inversions wrought by such conceptual shifts. Indeed, he asks us to appreciate and "savour the radicalism" (Clark 2015a, p. 4). Still, the proposed revisions of which Clark speaks are still piecemeal; they do not constitute a wholesale replacement of previous thinking. There is a conservative streak in Clark's thinking because he, like many others, continues to assume that the notion of representational content must be retained. In contrast, RECtification does not simply challenge some very central assumptions of the classical way of thinking about cognition, it proposes uprooting that conception completely.

REC seeks to do much more than cast new light on psychotechtonics: it asks us to do more than rethink the basic architecture of mind or the staging and functioning of cognitive processes. It asks us to rethink our conception of the very nature of basic minds, abandoning altogether the idea that there are contentful mental states at the roots of cognition. In seeking to clarify the true character of cognition REC aims to promote truly revolutionary ways of thinking about mind and cognition.

Conceptual revolutions are rare, to be sure. Yet following the REC path seems *bona fide* revolutionary precisely because in doing so it to press for "the replacement of a whole system of concepts and rules by a new system" (Thagard 1992, p. 6). Drawing on Thagard, if we understand the relationship between kind concepts in our sciences of mind as exhibiting a tree-like hierarchy, it is clear that REC presses not merely for some local conceptual branch jumping within a single tree, but switching to a new tree altogether. In this respect the enactivist framework, under REC's auspices, has the main hallmarks of previous conceptual revolutions. Notably, the Copernican revolution also required a fundamental ontological rethink, requiring us to reclassify "the earth as a kind of planet, when previously it had been taken to be *sui generis*" (Thagard 1992,

p. 36). Likewise, "Darwin did not simply pick away at the creationist conceptual structure: he produced an elaborate alternative edifice that supplanted it as a whole" (Thagard 1992, p. 36).

In a similar spirit, REC aims not merely to adjust certain aspects of the classical cognitivist vision but to supplant that outlook entirely. For example, in doing away with the idea that content can be found at the basis of cognition, REC does away with the content/vehicle distinction and hence vehicles. REC is a complete game changer—one that opens the door for truly new thinking about the mind that can take us beyond business-as-usual cognitive science. And, like its predecessor conceptual revolutions, once radicalized, enactivism will almost certainly have many hard-to-predict scientific and practical ramifications down the line.

Some, like Shapiro (2014a) will doubt that abandoning classical thinking can take enactivism "to the next step" (p. 215). The jury is still out on that question. Yet all parties can agree that only a RECtified account of cognition would be a fundamentally different and genuinely alternative account of mind to that promoted by classical cognitivism. In this light, perhaps, Shapiro is right to suggest that, once RECtified, what we will be looking at will not be "a more ferocious breed" of enactivism, but "a different animal altogether" (p. 215).

5 Conclusion

Philosophical clarity is needed just to find the enactivist path of virtue, let alone to stay on it. Those who want to remain philosophically clean and clear are advised against a RECless cognitive science.

References

Akins K (1996) Of sensory systems and the 'aboutness' of mental states. J Philos 93(7):337–372

Araújo D, Davids KW, Hristovski R (2006) The ecological dynamics of decision making in sport. Psychol Sport Exerc 7:653–676

Barandiaran XE, Egbert M (2013) Norm-establishing and norm-following in autonomous agency. Artif Life 20(1):5–28

Barandiaran XE, Di Paolo E, Rohde M (2009) Defining agency. Individuality, normativity, assymetry and spatiotemporality in action. J Adapt Behav 17(4):367–386

Beer RD (2009) Beyond control: the dynamics of brain–body–environment interaction in motor systems. In: Sternad D (ed) Progress in motor control: a multidisciplinary perspective. Springer, New York, pp 7–24

Bruineberg J, Rietveld E (2014) Self-organization, free energy minimization, and optimal grip on a field of affordances. Front Hum Neurosci. doi:10.3389/fnhum.2014.00599

Burge T (2010) The origins of objectivity. Oxford University Press, Oxford

Busemeyer JR, Townsend JT (1993) Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychol Rev 100:432–459

Byrge L, Sporns O, Smith LB (2014) Developmental process emerges from extended brain–body–behavior networks. Trends Cogn Sci 18(8):395–403. doi:10.1016/j.tics.2014.04.010

Cadez E, Heit E (2011) Forgetting curves emerge from dynamics of integrated memory. In: Proceedings of the thirty-third annual conference of the cognitive science society, pp 1673–1678

Cappuccio M, Froese T (eds) (2014) Enactive cognition at the edge of sense-making. Palgrave Macmillan, London

Chemero A (2009) Radical embodied cognitive science. MIT Press, Cambridge

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

Clark A (2013a) Expecting the world: perception, prediction and the origins of human knowledge. J Philos 110:469–496

Clark A (2013b) Whatever next? Predictive brains, situated agents and the future of cognitive science. Behav Brain Sci 36:181–204

Clark A (2015a) Embodied prediction. In: Metzinger T, Windt JM (eds) Open MIND: 25(T). MIND Group, Frankfurt am Main, pp 1–21. doi:10.15502/9783958570115

Clark A (2015b) Predicting peace: the end of the representation wars—a reply to Michael Madary. In: Metzinger T, Windt JM (eds) Open MIND: 25(T). MIND Group, Frankfurt am Main, pp 1–7. doi:10.15502/9783958570979

Clowes R, Mendonça D (2015) Representation redux: Is there still a useful role for representation to play in the context of embodied, dynamicist and situated theories of mind? New Ideas Psychol. doi:10.1016/j.newideapsych.2015.03.002

Colangelo A, Holden JG, Buchanon L, Van Orden GC (2004) Speculation about behavior, brain damage, and self-organization: the other way to herd a cat. Brain Lang 90:151–159

Cooke E, Myin E (2011) Is trilled smell possible? How the structure of olfaction determines the phenomenology of smell. J Conscious Stud 18(11–12):59–95

Dale R, Roche JM, Snyder K, McCall R (2008) Exploring action dynamics as an index of paired-associate learning. PLoS ONE 3:e1728. doi:10.1371/journal.pone.0001728

De Jaegher H, Di Paolo E (2007) Participatory sense-making: an enactive approach to social cognition. Phenomenol Cogn Sci 6(4):485–507

De Jesus P (2015) Autopoietic enactivism, phenomenology and the deep continuity between life and mind. Phenomenol Cogn Sci. doi:10.1007/s11097-015-9414-2

Di Paolo E (2005) Autopoiesis, adaptivity, teleology, agency. Phenomenol Cogn Sci 4:429–452

Dretske F (1988) Explaining behaviour: reasons in a world of causes. MIT Press, Cambridge

Elman JL (1995) Language as a dynamical system. In: Port RF, van Gelder T (eds) Mind as motion. MIT Press, Cambridge, pp 195–225

Fodor JA (1981) Representations: philosophical essays on the foundations of cognitive science. MIT Press, Cambridge

Fodor JA (1991) Fodor's guide to mental representation. In: Greenwood (ed) The future of folk psychology. Cambridge University Press, Cambridge

Friston KJ (2010) The free-energy principle: a unified brain theory? Nat Neurosci 11:127–138

Friston KJ, Stephan KE (2007) Free-energy and the brain. Synthese 159:417–458

Froese T, Di Paolo EA (2011) The enactive approach: theoretical sketches from cell to society. Pragmat Cogn 19(1):1–36. doi:10.1075/pc.19.1.01fro

Froese T, Ziemke T (2009) Enactive artificial intelligence: investigating the systemic organization of life and mind. Artif Intell 173(2009):466–500

Gladziejewski P (2015a) Explaining cognitive phenomena with internal representations: a mechanistic perspective. Stud Log Gramm Rhetor 40(53):63–90

Gładziejewski P (2015b) Predictive coding and representationalism. Synthese. doi:10.1007/s11229-015-0762-9

Godfrey-Smith P (2006) Mental representation, naturalism and teleosemantics. In: Macdonald G, Papineau D (eds) Teleosemantics. Oxford University Press, Oxford, pp 42–68

Grush R (1997) The architecture of representation. Philos Psychol 10:5–23

Grush R (2004) The emulation theory of representation: motor control, imagery and perception. Behav Brain Sci 27:377–442

Heras-Escribano M, Noble J, de Pinedo M (2014) Enactivism, action and normativity: a Wittgensteinian analysis. Adapt Behav. doi:10.1177/1059712314557364

Hohwy J (2013) The predictive mind. Oxford University Press, Oxford

Hohwy J (2014) The self-evidencing brain. Noûs. doi:10.1111/nous.12062

Holden JG, Van Orden GC, Turvey MT (2009) Dispersion of response times reveals cognitive dynamics. Psychol Rev 116:318–342

Hurley S (1998) Consciousness in action. Harvard University Press, Cambridge

Hurley S, Noë A (2003) Neural plasticity and consciousness. Biol Philos 18:131–168

Hutto DD (2005) Knowing what? Radical versus conservative enactivism. Phenomenol Cogn Sci 4:389–405

Hutto DD (2011) Enactivism: Why be radical? In: Bredekamp H, Krois JM (eds) Sehen und Handeln. Akademie Verlag, Berlin, pp 21–44

Hutto DD (2014) Tractatus Logico-Philosophicus, untimely review. Ludwig Wittgenstein. An International Review of Philosophy, Topoi. doi:10.1007/s11245-014-9291-2

Hutto DD, Myin E (2013) Radicalizing enactivism. MIT Press, Cambridge

Hutto DD, Satne G (2015) The natural origins of content. Philosophia 43(3). doi:10.1007/s11406-015-9644-0

Jonas H (1968) Biological foundations of individuality. Int Philos Q 8(2):231–251

Kant I (1790/1987) Critique of judgment (trans: Pluhar WS). Hacket Publishing Company, Indianapolis, IN

Kelso JAS (1995) Dynamic patterns. MIT Press, Cambridge

Kiverstein J, Rietveld E (2015) The primacy of skilled intentionality: on Hutto & Satne's the natural origins of content. Philosophia 43(3). doi:10.1007/s11406-015-9645-z

Malafouris M (2013) How things shape the mind: a theory of material engagement. MIT Press, Cambridge, MA

Maturana H, Varela F (1980) Autopoiesis and cognition. Reidel, Dordrecht

McKinstry C, Dale R, Spivey MJ (2008) Action dynamics reveal parallel competition in decision making. Psychol Sci 19:22–24

Menary R (2007) Cognitive integration: mind and cognition unbounded. Palgrave, Basingstoke

Menary R (2015) Mathematical cognition—a case of enculturation. In: Metzinger T, Windt JM (eds) Open MIND: 25(T). MIND Group, Frankfurt am Main. doi:10.15502/9783958570818

Menary R (forthcoming) What? Now. Predictive coding and enculturation. In: Metzinger T, Windt JM (eds) Open MIND: 25(T). MIND Group, Frankfurt am Main

Millikan R (1993) White queen psychology and other essays for Alice. MIT Press, Cambridge

Millikan RG (2005) Language: a biological model. Oxford University Press, Oxford

Myin E, Hutto DD (2015) REC: just radical enough. Stud Log Gramm Rhetor 41(54):61–71. doi:10.1515/slgr-2015-0020

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

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

Noë A (2012) Varieties of presence. Harvard University Press, Cambridge

O'Brien G, Opie J (2004) Notes toward a structuralist theory of mental representation. In: Clapin H, Staines P, Slezak P (eds) Representation in mind: new approaches to mental representation. Elsevier, Oxford, pp 1–20

O'Brien G, Opie J (2015) Intentionality lite or analog content? a response to Hutto and Satne. Philosophia 43(3). doi:10.1007/s11406-015-9623-5

O'Regan K (2011) Why red doesn't sound like a bell. Explaining the feel of consciousness. Oxford University Press, Oxford

O'Regan JK, Noë A (2001) A sensorimotor account of vision and visual consciousness. Behav Brain Sci 24:939–1031

Orlandi N (2014) The innocent eye: why vision is not a cognitive process. Oxford University Press, Oxford

Phattanasri P, Chiel HJ, Beer RD (2007) The dynamics of associative learning in evolved model circuits. Adapt Behav 15:377–396

Pietroski P (1992) Intentionality and teleological error. Pac Philos Q 73:267–282

Putnam H (1988) Representation and reality. MIT Press, Cambridge

Ramsey WM (2007) Representation reconsidered. Cambridge University Press, Cambridge

Rey G (2015) Representation. In: Garvey J (ed) The Bloomsbury companion to philosophy of mind. Bloomsbury, London

Rietveld E (2008) Situated normativity: the normative aspect of embodied cognition in unreflective action. Mind 117(468):973–1001

Rietveld E, Kiverstein J (2014) A rich landscape of affordances. Ecol Psychol 26(4):325–352

Riley MA, Holden JG (2012) Dynamics of cognition. Wiley Interdiscip Rev Cogn Sci 3(6):593–606. doi:10.1002/wcs.1200

Rueckl J (2002) The dynamics of visual word recognition. Ecol Psychol 14:5–19

Rupert R (2011) Embodiment, consciousness, and the massively representational mind. Philos Top 39(1):99–120

Shapiro L (2014a). Radicalizing enactivism: basic minds without content (Review). Mind 123(489):213–220

Shapiro L (2014b) When is cognition embodied? In: Kriegel U (ed) Current controversies in philosophy of mind. Routledge, London

Shea N (2013) Naturalising representational content. Philos Compass 8(5):496–509

Stephen DF, Dixon JA (2009) The self-organization of insight: entropy and power laws in problem solving. J Problem Solv 2:72–101

Stewart J, Gapenne O, Di Paolo E (eds) (2010) Enaction: toward a new paradigm for cognitive science. The MIT Press, Cambridge

Tabor W (2002) The value of symbolic computation. Ecol Psychol 14:21–52

Thagard P (1992) Conceptual revolutions. Princeton University Press, New Jersey

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

Travis C (2004) The silence of the senses. Mind 113(449):57–94

Turvey MT (1990) Coordination. Am Psychol 45:938–953

Van Dijk L, Withagen R, Bongers RM (2015) Information without content: a Gibsonian reply to enactivists' worries. Cognition 134:210–214

Wijnants ML, Hasselman F, Cox RFA, Bosman AMT, Van Orden GC (2012a) An interaction-dominant perspective on reading fluency and dyslexia. Ann Dyslexia 62:100–119

Wijnants ML, Cox RFA, Hasselman F, Bosman AMT, Van Orden GC (2012b) A trade-off study revealing nested timescales of constraint. Front Fractal Physiol 3:116. doi:10.3389/fphys.2012.00116

Wittgenstein L (1953) Philosophical investigations (Anscombe GEM Trans), 3rd edn. Basil Blackwell, Oxford

Zednik C (2011) The nature of dynamical explanation. Philos Sci 78:238–263