Radically Enactive Cognition in Our Grasp

Daniel D. Hutto

University of Hertfordshire

"the question is not to follow out a more or less valid theory

but to build with whatever materials are at hand.

The inevitable must be accepted and turned to advantage"

- Napoleon Bonaparte

Abstract

Radically Embodied/Enactive accounts of Cognition, REC, propose to fundamentally shift the way cognitive scientists think about the basic nature of mentality. This paper argues that focusing on the sophisticated but unplanned character of human manual activity enables such accounts to address a standard worry about their scope and reach. A counter proposal for handling such cases by defenders of Conservative Embodied/Enactive account of Cognition, CEC, is examined and found wanting. CEC accounts make appeal to Action Oriented Representations (AORs) to do the work that fans of REC argue is done without representational mediation. It is argued that naturalistically inclined defenders of CEC face a crippling dilemma.

1. Reckoning with REC

For those working in the sciences of the mind these are interesting times. Revolution is, yet again, in the air. This time it has come in the form of new wave thinking about the basic nature of mind of the sort associated with radically embodied or enactive approaches to cognition; REC for short. REC approaches are marked out by their uncompromising and thoroughgoing rejection of intellectualism about the basic nature of mentality. As Varela, Thompson and Rosch (1991) saw it the defining characteristic of this movement is its opposition to those theories of mind that “take representation as their central notion” (p. 172). The most central and important negative claim of REC is its denial that all forms of mental activity depend on the construction of internal models of worldly properties and states of affairs by means of representing its various features on the basis of retrieved information.

Not since the ousting of behaviourism with the advent of the most recent cognitive revolution has there been such a root and branch challenge to widely accepted assumptions about the very nature of mentality. In a remarkable reversal of fortune, it is now a live question to what extent, *if any*, representational and computational theories of the mind – those that have dominated for so long – ought to play a fundamental role in our explanatory framework for understanding intelligent activity. Defenders of REC approaches argue that representation and computation are neither definitive of, nor provide the basis of, *all* mentality.

From the side-lines, interested onlookers might be forgiven for thinking the revolution is already over; embodied and enactive ways of thinking are already comfortably ensconced, having established deep roots in a number of disciplines. Far from merely being at the gates, the Barbarians are, it seems, now occupying the local cafés and wine bars in the heart of the city. Even those who most regret this development are prepared to acknowledge that, *de facto*, there has been a major sea change. Lamenting the rise of a pragmatist trend in cognitive science, Fodor (2008) acknowledges that REC-style thinking in cognitive science is now the 'mainstream'. He puts this down to an infectious disease of thought ('a bad cold' – as he puts it p. 10). Others are edgily aware of the spectre of REC approaches "haunting the laboratories of cognitive science" (Goldman and de Vignemont 2009, p. 154). 'Pervasive and unwelcome' is the verdict of these authors: REC may be everywhere but is something to be cured or exorcised, as soon as possible.

Despite their growing popularity, which some hope is nothing more than a short-lived trend, REC approaches remain hotly contested. Certainly, it is true that there has yet to be a definitive articulation of the core and unifying assumptions of embodied and enactive approaches to cognition – EC approaches – radical or otherwise. Indeed, there is some reason to doubt that it will be possible to group together all of the offerings that currently travel under the banner of EC by identifying their commitment to a set of well-defined core theoretical tenets (see Shapiro 2011, p. 3). Nevertheless, if REC approaches, in particular, are to maintain credibility and avoid charges of simply riding the crest of a fashionable wave, at a bare minimum, serious objections from the old guard should be convincingly answered.

Some criticisms are easier to deal with than others. One line of argument draws on observations about the proper order and requirements of cognitive explanations. Fodor (2008) hopes to dispatch REC with one fell blow by observing that positing of representationally-based thinking is the minimum requirement for explaining any and all activity that deserves the accolade ‘intelligent responding’ – an observation predicated on the assumption that we can draw a principled bright line between what is properly cognitive and what is not. Accordingly, he insists that we have no choice but to accept that “the ability to think the kind of thoughts that have truth-values is, in the nature of the case, prior to the ability to plan a course of action. The reason is perfectly transparent: Acting on plans (as opposed to, say, merely behaving reflexively or just thrashing about) requires being able to think about the world” (p. 13).

In a nutshell, this is Fodor’s master argument for thinking that pragmatist REC-style approaches to the mind *must* be false: for to think the kinds of thoughts that have truth-values is to think thoughts with representational content and, presumably, to make plans requires manipulating these representations (and their components) computationally.[1] So, in short, if *all* bona fide intelligent action involves planning, and all *bona fide* planning involves computing and representation then this is bad news from the frontline for REC rebels.

Without a doubt some problems, indeed, perhaps whole classes of problems, are best addressed through advanced careful planning – planning of the sort that requires the rule-governed manipulation of truth-evaluable representations. Sometimes it is not only advisable, but utterly necessary to stand back and assess a situation in a relatively detached manner, drawing explicitly on general background propositional knowledge of situations of a similar type and using that knowledge to decide – say, by means of deduction and inference – what would be the correct or most effective approach. This can be done before initiating any action or receiving any live feedback from the environment. This is the preferred strategy for dealing with situations, such as defusing bombs, in which a trial and error approach is not advisable. Making use of remote representations also works equally even for more mundane types of tasks – those that include, for example, figuring out the best route from the train station to one's hotel in a foreign city where one seeks to do this in advance from the comfort of one's office, long before boarding a plane.

Such intelligent pre-planning can be done at a remove, and reliably, if it is possible to exploit and manipulate symbolic representations of the target domain on the assumption that one has the requisite background knowledge and can bring that knowledge to bear. This will work if the domain itself is stable over time since that will ensure that any stored representations remain up to date and accurate. By using representations of a well-behaved domain's features and properties, and having a means of knowing, determinately, what to expect of it if it changes under specific modifications and permutations, it is possible for a problem solver to plan how to act within it without ever having to (or indeed ever having had to) interact with it in a first-hand manner. This is, of course, the ideal end state of high theoretical science.

As linguistic beings, humans are representation mongers of this sort and thus regularly adopt this basic strategy to solve problems. Our cultural heritage provides us with a store of represented knowledge – in many and various formats – that enables us to do so, successfully, under the sorts of conditions just mentioned. But it hardly follows that this type of cognitive engagement is the basis of, required for, or indeed suitable for, all sorts of tasks – *always and everywhere*. Echoing Ryle (1949), Noë (2009) hits the nail on the head, noting that, "the real problem with the intellectualist picture is that it takes rational deliberation to be *the most basic kind of cognitive operation*" (p. 99, emphasis added).

Intellectualism of this unadulterated kind – of the sort that assumes the existence of strongly detached symbolic representations of target domains – has fallen on hard times. It finds only a few hard-core adherents in today's cognitive sciences. Indeed, if anything has promoted the fortunes of REC it has been the dismal failure of this sort of rules and representation approach when it comes to dealing with *the most basic forms* of intelligent activity. This is the headline-grabbing lesson of recent efforts in robotics and artificial intelligence, which have provided a series of existence proofs against strong representationalism about basic cognition.

Pioneering work by Brooks (1991a, 1991b), for example, reveals that intellectualism is a bad starting point when thinking about how to build robots that actually work. There are important lessons to learn by paying attention to architectonic requirements of robots that are able to complete quite basic sorts of tasks, such as navigating rooms while avoiding objects or recognizing simple geometrical forms and shapes. Inverting standard intellectualist thinking, Brooks famously rejected the Sense-Model-Plan-Act approach, and built robots that dynamically and frequently sample features of their local environments in order to directly guide their responses, rather than going through the extra steps of generating and working with *descriptions* of those environments. These first generation behaviour-based robots, and those that followed after them, succeed precisely because the robots' behaviours are guided by continuous, temporally extended interactions with aspects of their environments rather than working on the basis of represented internal knowledge about those domains, knowledge that would presumably be stored somewhere wholly in the robots' innards. The guiding principle behind Brooks' so-called subsumption architectures is that sensing is *directly connected* with appropriate responding without *representational* mediation. Crucially, the great success of these artificial agents demonstrates that it is possible for a being to act intelligently without creating and relying on internal representation and models. Very much in line with theoretical worries raised by the frame problem, it may even be that, when it comes to basic cognition, this is the only real possibility.

Not just artifice but nature too provides additional support for the same conclusion. Cricket phonotaxis (Webb 1994, 1996) is a vivid example of how successful on-line and successful navigation takes place in the wild, apparently without the need for representations or their manipulation. Female crickets locate mates by attending to the first notes of male songs, frequently adjusting the path of their approach accordingly. They only manage this because the male songs that they attune to have a unique pattern and rhythm – one that suits the particular activation profiles of the female interneurons. The capacity of these animals to adjust their behaviour when successfully locating mates requires them to engage in a continuous interactive process of engagement with the environment. In doing so they exploit special features of their non-neural bodies – including the unique design of their auditory mechanism – as well as special features of the environment – the characteristic pattern of the male songs. In this case a beautiful cooperation arises because of the way the cricket's body and wider environment features enable successful navigational activity – activity that involves nothing more than a series of dynamic and regular embodied interactions.

For reasons of space I will not rehearse in full detail precisely how behaviour-based robots or insects make their way in the world. These cases are well known and much discussed (for excellent summaries in greater detail, including other examples see Wheeler 2005, ch. 8 and Shapiro 2011, ch. 5). For the purposes of this essay it suffices to note that when bolstered by the articulation of a supporting theoretical framework, one easily provided by dynamical systems theory, these observations offer a serious and well-known challenge to the representationalist assumptions of intellectualist cognitive science (Beer 1998, 2000, Thompson 2007, Garzón 2008, Chemero 2009).

In sum, what the foregoing reflections teach us is that there are cases in which bodily and environmental factors play ineliminable and non-trivial parts in making certain types of cognition possible. A familiar intellectualist response to these sorts of examples is to try to cast these wider contributions as playing no more than causal supporting roles that, even if necessary to enable cognition do not constitute or form part of it. For reasons that should be obvious from the foregoing discussion, it is not clear how one might motivate this interpretation and make it stick with respect to the sorts of cases just described.

In rejecting representationalism, REC takes at face value what attending to the architectonic details of how these agents work suggests – i.e. that the specified bodily and environmental factors are fully *equal partners* in constituting the embodied, enactive intelligence and cognition of these artificial and natural agents. Accordingly, although for certain practical purposes and interventions it may be necessary to carve off and focus on specific causally contributing factors in isolation, the cognitive activity itself cannot be seen as other than “a cyclical and dynamic process, with no nonarbitrary start, finish, or discrete steps” (Hurley 2008, p. 12, see also Garzón 2008, p. 388). Or, put otherwise, when it comes to understanding cognitive acts “the agent and the environment are non-linearly coupled, they, together constitute a nondecomposable system” (Chemero 2009, p. 386).

In promoting this sort of line, REC flags up the ‘real danger’ that “the explanatory utility of representation talk may evaporate altogether” (Wheeler 2005, p. 200). As Shapiro (2011) notes, the interesting question is whether an anti-representationalist paradigm has real prospects of *replacing* intellectualist cognitive science altogether. And, as indicated above, he is right to suppose that the two main, complementary “sources of support for Replacement come from (i) work that treats cognition as emerging from a dynamical system and (ii) studies of autonomous robots” (p. 115). While this is a potentially powerful cocktail, it remains to be seen just how far it might take us. For to make a convincing case for their far-reaching revolutionary ambitions, proponents of REC must take the next step and “argue that *much or most cognition* can be built on the same principles that underlie the robot’s intelligence” (Shapiro 2011, p. 116, emphasis added).

Rather than denying that there can be no such thing as non-representational cognition, intellectualists might take heart from this challenge and agree to split the difference, allowing that very basic forms of cognition – of the sort exemplified by robot and insect intelligence – might be suitable for REC treatment but not the rest. This is to adopt a kind of containment strategy – a kind of theoretical kettling or corralling. Intellectualists might be tempted to concede that supporters of the radical left have a point, up to a point –

allowing that “representations are not necessary for coordinated responses to an environment with which one is dynamically engaged” (Shapiro 2011, p. 153). But this concession would be made in the secure knowledge that it “would support only the conclusion that agents do not require representations for certain kind of activities. However, a stronger conclusion, for instance that cognition never requires representational states, does not follow” (Shapiro 2011, p. 153).

REC approaches would be, accordingly, of limited value on the assumption that they won’t scale up. Call this the Scope objection. It allows one to accept certain lessons learned from the lab and nature while safe in the knowledge that even if representations are not needed to explain the most basic forms of cognition that this in no way poses an interesting threat to intellectualism since the sorts of cases in question “represent too thin a slice of the full cognitive spectrum” (Shapiro 2011, p. 156). This is in line with the oft-cited claim that some behaviour is too off-line and representation hungry to be explained without appeal to the manipulation of symbolic representations. In particular, non-representational cognition, which might do for simple robots and animals, isn’t capable of explaining properly world-engaging, human forms of cognition. But should that assessment prove mistaken – if REC approaches were to make substantial in-roads in this latter domain – then the boot might just be on the other foot. For it might turn out that representationally hungry tasks only make up a very small portion of mental activity; representationally-based cognition might be just the tip of the cognitive iceberg.

2. A Helping Hand

This is where reflection on the special prowess of the human hand comes in handy. It cannot be denied that a great deal of human manual activity is connected with sophisticated forms of cognition.

Milner and Goodale's (1995) famous experiments reveal that humans can perform remarkably demanding manual acts, with precision – acts requiring the exercise of very fine-grained motor capacities, such as posting items through slots with changing orientations – even when they lack the capacity to explicitly report upon or describe visual scenes they are dealing with.

Nor, with only rare exceptions, is it credible that humans normally learn how to use their hands in these sorts of ways by means of explicit, representationally mediated instruction, the rules for which only later becoming submerged and tacit. It is not as if children are taught by their caregivers through explicit description how to grasp or reach for items. Far more plausibly, is the hypothesis that we become handy through a prolonged history of interactive encounters – through practice and habit. An individual's manual know how and skills are best explained entirely by appeal to a history of previous engagements and not by the acquisition of some set of internally stored mental rules and representations. To invoke the favourite poetic motto of enactivists, this looks, essentially, to be a process of 'laying down a path in walking' or in this case, 'handling'.

It is possible that the special manual abilities of humans are sophisticated enough to have provided the platform and spurred on other major cognitive developments. Some very strong claims have been made about the critical importance of the ways in which we use our hands in this regard – ways that some believe are responsible for enabling the emergence of other distinctively forms of human cognition, consciousness and culture. For example, Tallis (2003) regards our special brand of manual activity as the ultimate source of our awakening to self-consciousness. He tells us, “Herein lies the true genius of the hand: out of fractionated finger movements comes an infinite variety of grips and their combinations. And from this variety in turn comes choice – not only in what we do … but in how we do it … [and w]ith choice comes consciousness of acting” (p. 175).

If Tallis is to be believed, “Between the non-stereotyped prehensions of hominid hand and the stereotyped graspings of the animal paw there is opened a gap which requires, and so creates, the possibility of apprehension to cross it” (p. 36). These claims are tempered by the remark that “We may think of the emergence of distinctive capabilities of the human hand as *lighting a fuse on a long process* that entrained many other parts of the human body and many other faculties as it unfolded” (p. 6, emphasis added). This allows for the possibility that, “The crucially important differences between human and non-human hands do not alone account for the infinitely complex phenomenon, unique in the order of the universe, of human culture. It is not so much the differences – which are very important – but the ability to make much of the differences” (p. 33).[2]

Whatever is ultimately concluded about the defensibility of this last set of claims, the point is that if it should turn out that much human manual activity is best explained without appeal to the manipulation of rules or representations then defenders of REC will have made significant progress in addressing the scope challenge. REC approaches will have shown the capacity to advance well beyond dealing with the antics of a behaviour-based robots and insects, having moved deep in the heart of distinctively human cognition.

Are there further grounds for thinking that manual activity is best explained in a representation-free way? Tallis's (2003) philosophically astute and empirically informed examination of the hand provides an excellent starting point for addressing this question. He claims that "the hand [is] … an organ of cognition", and is so 'in its own right' (p. 28, p. 31). This is not to say that the hand works in isolation from the brain, indeed, Tallis stresses that the hand – for him, the tool of tools – is the "brain's most versatile and intelligent lieutenant" (p. 22). Of course, this way of putting things suggests that the hand is, when all goes well, in some way nothing but a faithful subordinate – one that works under top-down instruction and guidance from above. This underestimates the bi-directional interplay between manual and brain activity - interplay of the sort that explains why the distinctive manual dexterity of *Homo sapiens,* which sets us apart even from other primates who also have remarkable abilities in this regard, was likely one of the 'main drivers' of the growth of the human brain (p. 22).

These ideas can be taken much further if one fully rejects what Tallis calls the standard ploy.

> While it is perfectly obvious that voluntary activity must be built up out of involuntary mechanisms, there are *profound problems* in understanding this. There are particular problems with *the standard ploy* invoked by movement physiologists: proposing that

> the automation incorporates 'calculations' that the brain (or part of it) 'does', which permit customisation of the programs to the singularities of the individual action (p. 65, emphases added)

Invoking the standard ploy amounts to making a hand waving and anthropocentric appeal to representational contents so as to specify and fill out hypothesized motor plans and motor programmes that supply instructive orders from on high, lending intelligence to and directing manual activity. For example, on this view motor plans, intentions and programs are understood as "propositional attitudes with contents of the form 'let [my] effector E perform motor act M with respect to goal-object G'" (Goldman 2009, p. 238).[3] The trouble is that even if we imagine that such representational contents exist, it is difficult to see how they could do the required work. The only chance they could have of specifying what is to be done, and how it is to be done, would be if they go beyond issuing very general and abstract instructions of the sort that Goldman gestures at above.

Only very fine-grained instructions would be capable of directing or controlling specific acts of manual activity successfully. This raises a number of questions. How do brains decide which general kind of motor act, M, is the appropriate sort of motor act to use in the situation at hand? This, alone, is no simple business – given the incredible variety of possible manual acts.[4]

And, even if we put that concern aside, proponents of the view that brains can initiate and control manual acts by traditional intellectualist means are left with the problem of explaining how and on what basis brain decides how to execute any given act. A major problem for traditional forms of intellectualism is that the requirements for successfully performing any particular motor act are tied to a unique and changing context. For example, even if everything else remains static the speed, angle of approach and grip aperture need to be altered appropriately at successive stages as one does something as simple as picking up a coffee cup. In a nutshell:

> a particular challenge … has been to explain how cognition and perception processes regulate complex multi-articular actions in dynamic environments. The problem seeks to ascertain how the many degrees of freedom of the human motor system (roughly speaking the many component parts such as muscles, limb segments, and joints) can be regulated by an internally represented algorithm … and how the motor plan copes with the ongoing interaction between the motor system and energy fluxes surrounding the system, e.g., frictional forces and gravitational forces ... Not even the attempt to distinguish between the motor plan and the motor program has alleviated the problem in the literature (Summers and Anson, 2009) (Araújo & Davids 2011, p. 12).

Successful manual activity requires bespoke and on the fly customisation. Hence, it is deeply implausible that brain can simply identify what is required for successfully completing a certain type of activity and simply issuing general instructions to be carried out in form of ordering pre-programmed routines to be carried out. The implausibility of this suggestion is underscored by the fact that "most of the things we do are unique even though they may have stereotyped components" (Tallis 2003, p. 67). Not surprisingly, human manual activity – despite its unique complexities – seems to depend on interactions between the brain, body and environmental interactions which involve essentially the same kinds of dynamic interactive feedback and temporally extended engagements needed to explain the intelligent antics of behaviour-based robots and insects.

It must be noted that concomitant with abandoning the standard ploy in favour of a REC based approach comes the admission that what we are dealing in most cases of manual activity are not strictly speaking actions. This will surely be so if we operate with a strict concept of action – one that insists on a constitutive connection between actions and intentional states, where the latter are conceived of as requiring the existence of propositional attitudes of some sort. But all that follows from this, as Rowlands (2006) observes, is that "most of what we do does not count as action" (p. 97).

Respecting the stipulated criterion on what is required for action, many philosophers acknowledge the existence of non-intentional doings, motivated activities and/or deeds. For example, Velleman (2000) recognizes the need to:

> define a category of ungoverned activities, distinct from mere happenings, on the one hand, and from autonomous actions, on the other. This category contains the things one does rather than merely undergoes, but that one somehow fails to regulate in the manner that separates autonomous human action from merely motivated activity (p. 4).

On the face of it, the great bulk of animal doings takes the form of sophisticated forms of highly coordinated, motivated activity that falls well short of action if acting requires forming explicit, if non-conscious, intentions and deliberate planning, at any level.

Picking up on Fodor's earlier remark, far from being mere 'thrashings about' or 'reflexive behaviors', such unplanned engagements appear to be quite skillful, and sometimes even expert, dealings with the world. If REC has the right resources for explaining the wide class of such doings then it has the potential to explain quite a lot of what matters to us when it comes to understanding mind and cognition.

3. The Non-Standard Ploy: Representationalism Rescued?

Despite all that has been said in favour of REC, many will baulk at going so far. There are weaker, and much more conservative and conciliatory ways of taking on board what is best in embodied and enactive ideas without abandoning intellectualism in a wholesale way. For example, intellectualists can happily accept that various facts about embodiment are causally necessary in making certain types of intelligent responding possible and in shaping its character without this concession in any way threatening the idea that cognition is wholly constituted by representational facts or properties.

Trivially, it is clearly true that what a creature perceives depends on contingent facts about the nature of its sensory apparatus – thus bats, dolphins and rattlesnakes perceive the world differently and perceive different things because they are differently embodied. Moreover, no one denies that what and how we perceive causally depends on what we do – thus, it is only by moving my head and eyes in particular ways that certain things become visible and audible. Obviously, these truisms in no way threaten intellectualism.

Things can be taken further still without rocking the boat too much. A more daring thesis, one that several authors have lighted upon, is that extended bodily states and processes might – at least on occasion – serve as representational or information-carrying vehicles. As such, they can play unique computational roles in enabling some forms of cognition (Clark 2008b). Or, in the lingo of Goldman and de Vignemont (2009) perhaps those attracted to embodied and enactive accounts of cognition should be taken as claiming that some mental representations are encoded in essentially bodily formats. These renderings of what enactive and embodied accounts have to offer are conservative with respect to a commitment to representationalism. They are perfectly compatible with asserting that "the mind is essentially a thinking or representing thing" (Clark 2008, p. 149); or that "the manipulation and use of representations is the primary job of the mind" (Dretske 1995, p. xiv).

Without breaking faith with intellectualism, Conservative Embodied/Enactive Cognition, or CEC, still allows one to recognize "the profound contributions that embodiment and embedding make" (Clark 2008b, p. 45). For those who endorse only CEC and not REC the new developments in cognitive science, far from posing a threat to the existing paradigm, can be seen as supplying new tools or 'welcome accessories' of considerable potential value that could augment intellectualist accounts of the mind.

CEC-style thinking is best exemplified by a recent bid to save the representationalist baby from the embodied bathwater, by arguing for the existence of action-oriented representations, or AORs. According to Wheeler (2008), who has done more than most to promote this view, an action-oriented representation is one that is:

(i) action-specific (tailored to a particular behaviour and designed to represent the world in terms of specifications for possible actions);

(ii) egocentric (features bearer-relative content as epitomized by spatial maps in an egocentric co-ordinate system);

(iii) intrinsically context-dependent (the explicit representation of context is eschewed in favour of situated special-purpose adaptive couplings that implicitly define the context of activity in their basic operating principles) (see also Wheeler 2005, p. 199).

Believing in AORs is consistent with accepting the neural assumption – an assumption that pays homage to the intuition that neural states and processes have a special cognitive status. Those attracted to this assumption believe it should be respected because, even though non-neural factors can qualify as representational vehicles, as it turns out, in most cases they do not. As such, the great majority cognitive explanations only ever involve representations that are wholly brain-bound. This is so even in cases in which it is necessary to making appeal to extra-neural but non-representational causal factors in order to explain the particular way that some particular intelligent activity unfolds. By accepting this last caveat, defenders of CEC allow that the full explanation of a given bout of intelligent behaviour need not be strongly instructional in character in the way demanded by the standard ploy.

Wheeler (2005) highlights the core features of CEC-style thinking, illustrating the role of AORs by appeal to the architecture of a simple behavior-based robot created by Francheschini, Pichon and Blanes (1992).

> The robot has a primary visual system made up of a layer of elementary motion detectors (EMDs). Since these components are sensitive only to movement, the primary visual system is blind at rest. What happens, however, is that the EMD layer *uses* relative motion *information*, generated by the robot's own bodily motion during the previous movement in the sequence to build a temporary *snap map* of the detected obstacles, constructed using an egocentric coordinate system. Then, in an equally temporary *motor map, information* concerning the angular bearing of those detected obstacles is *fused with information* concerning the angular bearing of the light source (supplied by a supplementary visual system) and a directional heading for the next movement is generated (Wheeler 2005, p. 196, first, second, fifth and sixth emphases added).

Wheeler (2005) considers and dismisses a number of possible minimal criteria for being an AOR – including appeal to selectionist strategies and decoupleability. After careful review, he settles on the idea that what is necessary and sufficient to distinguish behaviour-based systems that operate with AORs from those that do not is that the former systems exhibit arbitrariness and homuncularity. A system exhibits arbitrariness just when the equivalence class of different inner elements is fixed "by their capacity, when organized and exploited in the right ways, to carry specific items of information or bodies of information about the world" (p. 218). A system is a homuncular just when (a) it can be compartmentalized into a set of hierarchically organized communicating modules, and (b) each of those modules performs a well-defined sub-task that contributes towards the collective achievement of the overall adaptive solution.

For Wheeler, the linchpin holding this account of AORs together is that some cognitive systems are information processing systems. Thus:

> The connection between our two architectural features becomes clear once one learns that, in a homuncular analysis, the *communicating sub-systems* are conceptualized as *trafficking in the information that the inner vehicles carry*. So certain subsystems are interpreted as producing information that is then consumed downstream by other subsystems (p. 218, emphases added).

We can legitimately describe a cognitive system as employing AORs just in case it is a genuine source of adaptive richness and flexibility and it turns out that its subsystems use "information-bearing elements to stand in for worldly states of affairs in their communicative dealings" (Wheeler 2005, p. 219). Satisfaction of the above conditions is all that is required for the existence of weak or minimal representations. In the end, all of the weight in this account is placed on the idea that it suffices for minimal representations to be present in a system, S, if it manipulates and makes use of informational content in well-defined ways.

This minimal notion of representation is, no doubt, attractive to cognitive scientists. For anyone in the field it is utterly textbook to be told that information is a kind of basic commodity – the raw material of cognition. After all, "minds are basically processors of information; cognitive devices [for] receiving, storing, retrieving, modifying and transmitting information of various kinds" (Branquinho 2001, xii-xiii).

There is great latitude for thinking about the processes that enable this. Thus:

> Mental representations might come in a wide variety of forms; there being no commitment in the claim itself to a specific kind of representation or to a particular sort of representational vehicle…mental representations might be thought of as images, schemas, symbols, models, icons, sentences, maps and so on (Branquinho 2001, xiv).

Accordingly, representations or representational vehicles "are items in the mind or brain of a given system that in some sense 'mirror', or are mapped onto, other items or sets of items … in the world" (Branquinho 2001, xiv). But what makes something into a vehicle, the essence of representing, is that they bear or possess content. Content is key. Thus:

> The whole thrust of cognitive science is that there are sub-personal contents and sub-personal operations that are truly cognitive in the sense that these operations can be properly *explained only in terms of these contents* (Seager 1999, p. 27, emphasis added).

Dietrich and Markman (2003) define representations as "any internal state that mediates or plays a mediating role between a system's inputs and outputs *in virtue of that state's semantic content. We define semantic content in terms of information* causally responsible for the state, and in terms of the use to which that information is put" (p. 97, emphasis added). In sum, believing in AORs only requires acceptance of "the general idea of inner states that bear contents" (Clark 2002, p. 386).

Given this, it might be thought that accepting that at least some cognitive systems employ AORs is a no-brainer. Certainly, it seems that this must be true of human manual activity of the sort described in Section 2 – activity of the sort that, if my argument goes through, would enable fans of REC to answer the Scope Objection. After all, even in cases in which that sort of activity is not supported by focused, conscious perception, "the motor activity of the hand – reaching, gripping and manipulation – cannot function in the absence of what is usually called 'sensory information'" (Tallis 2003, p. 27). Indeed, "the information the hand needs to support its manipulative function is most clearly evident in the first stage … in reaching out prior to grasping, shaping, etc. Here the hand is under primarily visual control: the target is located, the relationship to the body determined, the motion initiated to home in on the target – these are all regulated [or, better, assisted] by sight, which measures what needs to be done and the progress of the doing" (p. 27).

With this in mind, it looks like manual activity is surely dependent on information processing activity of the sort that would qualify as involving AORs – hence it is the sort of activity better suited to CEC rather than a REC treatment. If so, any ground gained by supporters of REC in Section 2 would be lost.

4. The Hard Problem of Content

Despite some obvious attractions, the AOR story and the support it lends to CEC over REC, isn't beyond question. Indeed, as this concluding section argues anyone who favours CEC must face up to the hard problem of content – and I suggest that REC is a small price to pay to allow one to avoid that problem.

Before turning to that issue, it is worth highlighting an immediate concern about CEC, and its reliance on AORs. Appeal to AORs seems to secure the fate of minimal representations – winning a key metaphysical battle – only at the cost of losing a wider explanatory war. For, on the assumption that the AORs need not be decoupled in order to qualify as representations – an assumption Wheeler explicitly defends (2005, 2008), and with good reason (see Chemero 2009, ch. 3), defenders of CEC face the charge that "talk of representations in coupled systems may be too cheap, or too arbitrary, and thus adds little or nothing to an explanation of how these systems work" (Shapiro 2011, p. 147).

Chemero (2009) too voices this exact worry, noting that:

> the representational description of the system does not add much to our understanding of the system … [thus] despite the fact that one can cook up a representational story once one has the dynamical explanation, the representational gloss does not predict anything about the system's behaviour that could not be predicted by dynamical explanation alone (p. 77).

Although initially cast as a purely explanatory concern it is clear that this issue cannot be kept wholly free of metaphysical considerations. For instance, Chemero goes on to note that "in terms of the physics of the situation, the ball, the outfielder, and the intervening medium are just one connected thing. In effective tracking, that is, the outfielder, the ball, and the light reflected from the ball form a single coupled system. No explanatory purchase is gained by invoking representation here: in effective tracking, any internal parts of the agent that one might call mental representations are causally coupled with their targets" (Chemero 2009, p. 114).

Part of the trouble here is that there does not appear to be any clean cut way to decide, with precision, which systems actually satisfy the relevant conditions for being minimally representational systems. For example, there are diverging opinions about whether Watt's much discussed centrifugal governor – a device originally designed to ensure a constant operating speed in rotative steam engines – qualifies as a representational device. This is despite the fact that the relevant parties in the debate are fully agreed about the characteristics of the governor's internal design, which are quite elegant and simple. The positions of the device's spindle arms interact with and modifies the state of a valve which controls the engine's speed – when the arm is high the valve slows the engine, when the arm is low the valve increases engine speed.

In line with the criteria laid out in the previous section, Chemero (2009) concludes that:

> It is possible to view the governor's arms as [noncomputational] representations … It is the function of particular arm angles to change the state of the valve (the representation consumer) and so adapt it to the need to speed up or slow down. The governor was so designed … to play that role … it is both a *map* and a *controller*. *It is an action oriented representation*, standing for the current need to increase or decrease speed (p. 71, emphases added).[5]

Shapiro (2011) dissents. After careful review of this issue, he concludes that, “Watt’s design of the centrifugal governor *requires that* the angle of the flyball arms carry information about the flywheel speed in order to regulate the speed. Still, the function of the flyball arms is not to carry information, but rather to play a role in regulating the speed of the flywheel” (p. 147).

The point is that in order to carry out its control function the spindle arms must co-vary with the relevant changes in engine speed. That they carry information in this sense is an unavoidable, collateral feature of their design that enables them to perform their regulatory work. Apparently, this is not sufficient to qualify as being a true information processor in the relevant sense – hence the governor is misdescribed as making use of AORs.

Shapiro (2011) goes on to contrast the governor with other kinds of information using devices. For he holds, “Some devices surely do include components that have the function to carry information. The thermostat … is such a device. Thermostats contain a bimetal strip because the behavior of this strip carries information about [i.e. covaries with] temperature, and the room” (Shapiro 2011, p. 148).

The important difference is that although the governor’s arms carry information about the flywheel arms, the governor does not use that information in order to perform its control tasks. This is meant to mark the subtle but utterly critical difference between merely complex systems and properly cognitive complex systems. Apparently, the thermostatic systems are designed to use information about the temperature *as such* in carrying out their work.

Put as starkly as this, one might be forgiven for failing to see what changes so dramatically when it comes to the operation of thermostats as opposed to Watt governors. A thermostat regulates the system's temperature, maintaining it at a desired point. Its mechanisms exploit the properties of the bimetallic strip that – when all is well – responds in reliable ways to temperature changes, bending one way if heated and the opposite way if cooled.

The important difference between the two types of systems is not that there are more mechanisms or steps involved in regulating temperature by this means. Rather, the crucial distinction is meant to be that the bimetallic strips in thermostats have the systemic function of indicating specific desired temperatures to other subsystems that use those indications to regulate their behaviour. It is because they function in this special way that devices of this general type are representational – they exploit pre-existing indication relations giving them the function to indicate how things stand externally and use those indications in particular ways.[6] Following Dretske (1995) if such devices were naturally occurring then they would "have a content, something *they say or mean*, that does not depend on the existence of our purposes and intentions … [They would] have original intentionality, something they represent, say, or mean, that they do not get from us" (p. 8, emphasis added).

To qualify as representational, an inner state must play a special kind of role in a larger cognitive economy. Crudely, it must, so to speak, have the function of *saying or indicating* that things stand *thus and so*, and to be consumed by other systems because it says or indicates in that way. Only then will an internal state or structure meet Ramsey's (2007) job description challenge.

It is plausible that many of the states (or ensembles of states) of systems that enable basic cognition are merely (1) reliably caused by (or nomically depend upon) the occurrence of certain external features, and (2) disposed to produce certain effects (under specific conditions), and (3) have been selected because of their propensities for (1) and (2). Yet states or structures that only possess properties 1-3 fail to meet the job description challenge. They fail to qualify as truly representational mental states having the proper function of saying 'things stand thus and so', rather they – like the Watt governor – only have the proper function of guiding a system's responses with respect to specific kinds of worldly offerings.

Exactly, what else is required to be a representation-using system? Wheeler (2005) speaks of the need for communicative transactions between homuncular subsystems. This informational dealing is the basis of true cognition – nonetheless, he stresses that this does not imply that the sub-systems "in any literal sense understand that information" (p. 218).

Fair enough, but even if they literally lack understanding, it might be thought that at least such subsystems must be literally trading in informational content – using and fusing it – even if they don't understand what it says. But talk of using and fusing contents, although quite common, cannot be taken literally either. For it is not as if informational content is a kind of commodity that gets moved about and modified in various ways; information is not "like a parcel in the mail" (Shapiro 2011, p. 35).

This being so it seems that *bona fide* cognitive systems are not special because they *literally* use and manipulate informational content (not even content that they don't understand). They are special, because it is their function to convey informational content without actually manipulating it as such.

We are now getting down to brass tacks. For this story to work – there must at least be content that these subsystems have the special function to convey – there must be something that it is their function to say even if they don't understand what they are saying or what is said. But exactly what is informational content?

Dretske (1981) speaks about informational content as "the what-it-is-we-can-learn from a signal or message in contrast to how-much-we-can-learn" (p. 47). He makes clear that he understands a signal's informational content to be a kind of propositional content of the *de re* variety. Propositions or propositional contents have special properties – minimally, they are bearers of truth. Assuming that informational contents are propositional is presumably what allows Dretske to hold that when signals carry information to the senses they tell "us truly about another state of affairs" (p. 44).

This is, of course, quite compatible with holding that informational content lacks fully-fledged representational properties. Thus one can hold that informational content is supplied by the senses, which is not representational content, and that more is required for informational content to be properly representational.

It is at this juncture that defenders of AORs and CEC – at least those who subscribe to an explanatory naturalism – face a dilemma. Since so much hangs on this it is worth going very slowly over some familiar ground. In the opening passage of Dretske's *Knowledge and the Flow of Information* (1981) we find the *locus classicus* and foundational statement on how to understand information processing systems in a way that is both required by the CEC story and that expresses a commitment to explanatory naturalism.

> In the beginning there was information. The word came later ... information (though not meaning) [is] an *objective commodity*, something whose generation, transmission and reception do not require or in any way presuppose interpretative processes. One is therefore given a framework for understanding how meaning can evolve, how genuine cognitive systems – those with the resources for interpreting signals, holding beliefs, acquiring knowledge – can develop out of *lower-order, purely physical, information-processing mechanisms* ... Meaning, and the constellation of mental attitudes that exhibit it, are manufactured products. The raw material is information (p. vii, emphases added).

Any explanatory version of naturalism seeks to satisfy what Wheeler (2005) charmingly calls the Muggle constraint: "One's explanation of some phenomenon meets the Muggle constraint just when it appeals only to entities, states and processes that are wholly nonmagical in character. In other words, no spooky stuff" (p. 5).

It is widely supposed that the informational theory of content comfortably meet this constraint. At least, its defenders have attempted to convince us that, when promoting it, there is nothing up their sleeves. This is because, as Jacob (1997) emphasizes:

> *the relevant notion of information at stake in informational semantics is the notion involved in many areas of scientific investigation* as when it is said that a footprint or a

> fingerprint carries information about the individual whose footprint or fingerprint it is. In this sense, it may also be said that a fossil carries information about a past organism. The number of tree rings in a tree trunk carries information about the age of the tree (p. 45, emphasis added).

This picks out the relevant notion by means of examples. We can call it the notion of information-as-covariance. Although theorists quibble about the strength and scope of the degree of covariance required in order for informational relations to exist, there is consensus that s's being F 'carries information about' t's being H iff the occurrence of these states of affairs lawfully, or reliably enough, covary.

But here's the rub. Anything that deserves to be called content has special properties – e.g. truth, reference, implication – that make it logically distinct from and irreducible to mere covariance relations holding between states of affairs. While the latter notion is surely scientifically respectable, it isn't able to do the required work of explaining content. Put otherwise, if information is nothing but covariance then it is not any kind of content – at least not if content of the sort defined in terms of its truth bearing properties. The number of a tree's rings can covary with its age; this does not entail that the first state of affairs says or conveys anything true about the second, nor *vice versa*. The same goes for states that happen to be inside agents and which reliably correspond with external states of affairs – these too, in and of themselves, do not 'say' or 'mean' anything in virtue of instantiating covariance relations. Quite generally, covariation in and of itself neither suffices for, nor otherwise constitutes content, where content minimally requires the existence of truth bearing properties. Call this the Covariance doesn't Constitute/Confer Content (or CCC) principle.

The CCC undermines the assumption that covariation is the worldly source of informational content. There is no doubt the idea of information-as-covariance is widely used in sciences; hence, it is not a hostage to fortune for explanatory naturalists. But if the CCC is true, there is a gaping explanatory hole in the official story propounded by those who follow Dretske's lead. Anyone peddling such an account is surely violating the Muggle constraint and ought to be brought to the attention of the Ministry of Magic.[7]

One might opt for the first horn of this dilemma and retain the scientifically respectable notion of information-as-covariance, and thus retain one's naturalistic credentials while relinquishing the idea there is such a thing as informational content. That is the path I recommend, but it requires giving up on CEC since – as argued in the previous section – the minimal requirement for distinguishing informational *processing* systems is that they make use of AORs which are defined as content-bearing vehicles. But the distinction between vehicles and contents falls apart, at least at the relevant level, if there are no informational contents to bear.

To avoid this one might opt to be impaled on the second horn. This would be to accept that contentful properties exist even if they don't reduce to, or cannot be explained in terms of, covariance relations. If contentful properties and covariance properties are logically distinct they might still be systematically related. Hence, it might be hoped that contentful properties can be naturalistically explained by some other means (e.g. by some future physics). Alternatively, they could be posited as explanatory primitives – as metaphysical extras that might be externally related to covariance properties. Thus they might have the status that Chalmer's (2010) still assigns to qualia – they might require us to expand our understanding of the scope of the natural. Contentful properties might pick out properties that – like phenomenal properties – are irreducible to and exist alongside basic physical properties. If so, the explanatory project of naturalism with respect to them would look quite different – it would be to discover the set of fundamental bridging laws that explain how contentful properties relate to basic physical properties. That would be the only way to solve what we might call the hard problem of content.

Of course, one might try to avoid both horns by demonstrating the falsity of the CCC by showing how contentful properties – e.g. truth bearing properties – reduce to covariance properties (Good luck with that!). A more plausible dilemma-avoiding move would be to show that the notion of information that is in play in these accounts is, in fact, meatier than covariance but is nonetheless equally naturalistically respectable.

After all, Dretske talks of indication relations not covariance relations, though the two are often confused. Tellingly, in continuing the passage cited above Jacob (1997) remarks that "In all of these cases, it is not unreasonable to assume that the informational relation holds between an indicator and what it indicates (or a source) independently of the presence of an agent with propositional attitudes" (p. 45, emphasis added). In making this last point, he stresses that "the information or indication relation is going to be a relation between states or facts" (p. 49-50).

However, following Grice, Dretske is wont to think of indication as natural meaning – as in 'smoke' means 'fire'. But smoke *means* fire only if it indicates fire to someone. It makes no sense to talk of indication in the absence of a user. Indication is, at least, a three-place relation whereas covariance by contrast, is a two-place relation. To think of indication as the basis for informational semantics therefore is already to tacitly assume that there is more going on than mere covaration between states of affairs.[8] This raises questions about how exactly the notion of information-as-indication relates to its scientifically respectable cousin, the notion of information-as-covariance. Moreover, we might wonder if this notion has independent naturalistic credentials of its own. Until these questions are answered promoters of AORs and CEC – those that rely on the existence of informational content to distinguish genuine cognitive systems from all others – haven't really got off the starting blocks with their theorizing.

5. Epilogue: Decoding Information

The 'code' metaphor is rife in the cognitive sciences but the cost of taking it seriously is that one must face up to the Hard Problem of Content! In light of the problems with CEC-style stories highlighted above, we have reason to think that on-line sensory signals 'carry information' (in one sense) but not that they 'pass on' meaningful or contentful messages – contentful information that is used and fused to form inner representations. Unless we assume that pre-existing contents exist to be received through sensory contact the last thread of the analogy between basic cognitive systems and genuinely communicating systems breaks down at a crucial point.

In line with REC, there are alternative ways to understand cognitive activity as involving a complex series of systematic – but not contentfully mediated – interactions between well-tuned mechanisms (see, e.g., Hutto 2011a, 2011b; Hutto and Myin in preparation). RECers press for an understanding of basic mentality as literally constituted by, and to be understood in terms of, concrete patterns of environmentally situated organismic activity, nothing more or less. If they succeed, the above arguments should encourage more cautious CEC types – those trying to occupy the mid-left – to take a walk on this wild side.[9]

References

Araújo, D. and K. Davids (2011). "What Exactly is Acquired During Skill Acquisition?" *Journal of Consciousness Studies* 18(3-4): 7-23.

Branquinho, J. (2001). *The Foundations of Cognitive Science*. Oxford, Oxford University Press.

Brooks, R. (1991a). "New Approaches to Robotics." *Science* 253: 1227-1232.

Brooks, R. (1991b). "Intelligence without Representation." *Artificial Intelligence* 47: 139-159.

Chalmers, D. (2010) *The Character of Consciousness*. Oxford, Oxford University Press.

Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, MA, MIT Press.

Clark, A. (2002). "Skills, Spills and the Nature of Mindful Action." *Phenomenology and the Cognitive Sciences* 1: 385-387.

Clark, A. (2008a). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* Oxford, Oxford University Press.

Clark, A. (2008b). "Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind?" *Philosophy and Phenomenological Research* 76(1): 37-59.

Dietrich, E. and A. Markman (2003). "Discrete Thoughts: Why Cognition Must Use Discrete Representations." *Mind and Language* 18: 95-119.

Donald, M. (1999). Preconditions for the Evolution of Protolanguages. *The Descent of Mind: Psychological Perspectives on Hominid Evolution*. M. C. Corballis and S. E. G. Lea. Oxford, Oxford University Press.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA, MIT Press.

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

Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, MA, MIT Press.

Fodor, J. A. (2008). *LOT 2: The Language of Thought Revisited*. Oxford, Oxford University Press.

Francheschini, N., J.-M. Pichon, et al. (1992). "From Insect Vision to Robot Vision." *Philosophical Transactions of the Royal Society series B* (337): 283-294.

Gallagher, S. (2008). "Are Minimal Representations Still Representations?" *International Journal of Philosophical Studies*. 16(3): 351-369.

Gangopadhyay, N. (2011). "The Extended Mind: Born to Be Wild? A Lesson From Action Understanding." *Phenomenology and the Cognitive Sciences* 10(3): 377-397.

Garzón, F. C. (2008). "Towards a General Theory of Antirepresentationalism." *The British Journal for the Philosophy of Science* 59 (3): 259-292.

Goldman, A. I. (2009). "Mirroring, Simulating and Mindreading." *Mind and Language* 24(2): 235-252.

Goldman, A. I. and F. de Vignemont (2009). "Is Social Cognition Embodied?" *Trends in Cognitive Sciences* 13(4): 154-159.

Hurley, S. (2008). "The Shared Circuits Model: How Control, Mirroring and Simulation can Enable Imitation and Mindreading." *Behavioral and Brain Sciences* 27: 1-58.

Hutto, D. D. (2011). Enactivism: Why be Radical? *Sehen und Handeln*. H. Bredekamp and J. M. Krois. Berlin, Akademie Verlag: 21-44.

Hutto, D.D. (2011). “Philosophy of Mind’s New Lease on Life: Autopoietic Enactivism meets Teleosemiotics”. *Journal of Consciousness Studies* 18(5-6): 44-64.

Hutto D.D., Myin E. (in preparation). *Radicalizing Enactivism*. Cambridge, MA: MIT Press

Jacob, P. (1997). *What Minds Can Do*. Cambridge, Cambridge University Press.

Makin, G. (2000). *The Metaphysics of Meaning*. London, Routledge.

Milner, D. and M. Goodale (1995). *The Visual Brain in Action*. New York, Oxford University Press.

Noë, A. (2004). *Action in Perception*. Cambridge, MA, MIT Press.

Noë, A. (2009). *Out of Our Heads*. New York, Hill and Wang.

Ramsey, W. M. (2007). *Representation Reconsidered*. Cambridge, Cambridge University Press.

Rowlands, M. (2006). *Body Language*. Cambridge, MA, MIT Press.

Shapiro, L. (2011). *Embodied Cognition*. London, Routledge.

Summers, J. J. and G. A. Anson (2009). “Current Status of the Motor Programme Revisited.” *Human Movement Science* 28: 566-577.

Tallis, R. (2003). *The Hand: A Philosophical Inquiry into Human Being*. Edinburgh, Edinburgh University Press.

Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of the Mind*. Cambridge, MA, Harvard University Press.

Van Gelder, T. (1995). “What Might Cognition Be, If Not Computation.” *Journal of Philosophy* 92: 345-381.

Varela, F. J., E. Thompson, et al. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Velleman, J. D. (2000). *The Possibility of Practical Reason*. Oxford, Oxford University Press.

Webb, B. (1994). Robotic Experiments in Cricket Phototaxis. *From Animals to Animats 3: Proceedings of the Third Annual Conference on Simulation of Adaptive Behaviour*. D. Ciff, P. Husbands, J. A. Meyer and S. W. Wilson. Cambridge, MA, MIT Press.

Webb, B. (1996). “A Cricket Robot.” *Scientific American* 275(6): 62-67.

Wheeler, M. (2005). *Reconstructing the Cognitive World*. Cambridge, MA, MIT Press.

Wheeler, M. (2008). “Minimal Representing: A Response to Gallagher.” *International Journal of Philosophical Studies*. 16(3): 371-376.

[1] This intellectualist way of understanding the basic nature of minds taps into a long tradition stretching back at least as far as Plato; it was revived by Descartes in the modern era, and regained ascendency, most recently, through the work of Chomsky during the most recent cognitive revolution. As Noë (2009) observes:

"What these views have in common – and what they have bequeathed to cognitive science – is the idea that we are, in our truest nature, thinkers. It is this intellectualist background that shapes the way cognitive scientists think about human beings" (p. 98).

2 There are clear connections that might be forged between this view and Donald's (1999) claim that when it comes to understanding human cognition "the most critical element is a capacity for deliberately reviewing self-actions so as to experiment with them … It would be no exaggeration to say that this capacity is uniquely human, and forms the background for the whole of human culture, including language" (p. 142).

3 It is perhaps understandable that in seeking to make sense of this cognitive activity we are naturally inclined to assume the existence of representations that "include not only 'commands' and 'calculations', but also 'if-then' and other logical operations. This shows how it seems impossible to make sense of cerebral control – requisition and modification – of motor programs, to describe them in such a way that they deliver what is needed while avoiding anthropomorphisms" (Tallis 2003, p. 65). The problem is that "attributing to the brain, or parts of it, or neural circuits, the ability to do things that we, whole human beings, most certainly cannot do, seems unlikely to solve the puzzle" (Tallis 2003, p. 65)

4 At a pinch, one could give a short list of these, which could include: "grasping, seizing, pulling, plucking, picking, pinching, pressing, patting, poking, prodding, fumbling, squeezing, crushing, throttling, punching, rubbing, scratching, groping, stroking, caressing, fingering, drumming, shaping, lifting, flicking, catching, throwing, and much besides" (Tallis 2003, p. 22).

5 Notably, Chemero holds that the centrifugal governor is not a computer even though it can be regarded as a representational device and in the respect he does not break faith with the conclusion of Van Gelder's original analysis when he first introduced the example into the literature (Van Gelder 1995).

6 Thus "If we suppose that, through selection, an internal indicator acquired a biological function, the function to indicate something about the animal's surroundings, then we can say that this internal structure represents" (Dretske 1988, p. 94).

7 To make vivid what is at stake it is worth noting that early analytic philosophers were at home with the view that the world is ultimately and literally composed, at least in part, by 'propositions'. These were conceived of as bedrock Platonic entities – mentionable 'terms' which, when standing in the right complex relations, constitute judgeable objects of thought. In commenting on Russell's version of this idea, Makin underscores the features that parallel many of the properties that Dretske demands of informational content. He stresses that: "with

propositions, it is crucial to bear in mind that they are not, nor are they abstracted from, symbolic or linguistic or psychological entities … On the contrary, *they are conceived as fundamentally independent of both language and mind*. *Propositions are first and foremost the entities that enter into logical relations of implication, and hence also the primary bearers of truth* ... ‘truth'’ and 'implication' apply, in their primary senses, to propositions and only derivatively to the sentences expressing them” (Makin 2000, p. 11, emphasis added).

8 Others too have noticed this. For example, Ramsey (2007) comments on the peculiar features of the quasi-semantic indication relation as follows: “Dretske and many authors are somewhat unclear on the nature of this relation. While it is fairly clear what it means to say that state A nomically depends upon state B, it is much less clear how such a claim is supposed to translate into the claim that A is an indicator of B, or how we are to understand expressions like ‘information flow’ and ‘information carrying’” (p. 133).

9 I am not alone in trying to persuade those in the CEC brigade to make this shift, see also Gangopadhyay (2011).