

Tolerant Enactivist Cognitive Science

Thomas Raleigh

ABSTRACT: Enactivist (Embodied, Embedded, etc.) approaches in cognitive science and philosophy of mind are sometimes, though not always, conjoined with an anti-representational commitment. A weaker anti-representational claim is that ascribing representational content to internal/sub-personal processes is *not compulsory* when giving psychological explanations. A stronger anti-representational claim is that the very idea of ascribing representational content to internal/sub-personal processes is a *theoretical confusion*. This paper criticises some of the arguments made by Hutto & Myin (2013, 2017) for the stronger anti-representational claim and suggests that the default Enactivist view should be the weaker non-representational position.

1. Introduction

Over the past few decades, in both cognitive science and philosophy of mind, we have witnessed a very striking surge of interest in E-approaches – ‘Enactivist’, ‘Embodied’, ‘Embedded’, ‘Extended’. There is not any immediately obvious shared core-thesis which unites the various theories that are commonly grouped under these E-theoretical banners. Nor is there a single ancestral theory or paradigm from which they are all descended¹. However, if one had to identify a family resemblance which justifies the use of the ‘E-theoretic’ umbrella-term, it would presumably be that they all reject, in different ways and to different degrees, the ‘Good Old Fashioned’ computational paradigm, according to which cognition and intelligent behaviour should be understood in terms of computations defined over internal representational symbols *located inside* the cognizing subject².

A key issue then facing any E-theorist, an issue which has been used to taxonomise the various E-theorists into different camps, is the extent to which representational notions are retained or abandoned altogether when offering E-theoretic explanations. The label ‘radical’ was initially employed by Andy Clark (1997) to characterise the sort of strongly non-representational approach that he wished to *deny*:

‘Thesis of Radical Embodied Cognition: Structured, symbolic, representational and computational views of cognition are mistaken. Embodied cognition is best studied by means of noncomputational and nonrepresentational ideas and explanatory schemes...’ (Clark, 1997, 148)

This ‘radical’ appellation was subsequently adopted by some of the theorists who wished to *endorse* such a strongly non-representational approach and has been incorporated into the title of two recent monographs: Chemero’s ‘Radical Embodied Cognitive Science’ (2009) and

¹ Historical forebears that have been cited include: J. J. Gibson’s (1966, 1979) work on ecological optics, the ‘cybernetic’ theories of Wiener (1948) and Ashby (1956), the ethological work of Jacob Von Uexküll (1934 [2011]), W. Grey Walter’s (1950) early work in robotics. In philosophy, Merleau-Ponty (1962) is often mentioned.

² The locus classicus for the ‘good old fashioned’ approach in artificial intelligence is probably Newell and Simon (1961). Within psychology, Marr (1982) provides a canonical statement of the computational approach and applies it to the case of vision. In philosophy, Fodor (1975, 1980) is perhaps its most prominent defender.

Hutto & Myin's 'Radicalizing Enactivism' (2013)³. Following Hutto & Myin (2013, 2017) we can distinguish Radical E-approaches (REC), which (largely) eschew representational notions, from Conservative E-approaches (CEC), which retain representational notions alongside an E-theoretic emphasis on embodiment, enaction or environmental embedding etc.

To be clear, Radical E-theorists are not full-blown *eliminativists* about representational content – they stop short of denying the existence of representational mental states altogether. Hutto & Myin, for example, make clear that they only wish to deny that representational content is involved in '*basic cognition*':

'What is ruled out by REC is that content-involving mentality is basic and is found in any and all forms of mindedness... Like classic eliminativism, REC denies that basic mentality and cognition should be modelled in terms of propositional attitudes. But unlike classical eliminativism, REC does not claim that propositional attitude explanations are never appropriate. REC holds... that some organisms – language users, at least – are capable of genuinely contentful, representational modes of thinking and reasoning. ...REC happily concedes that some very important forms of cognition essentially depend on the interactions between propositional attitudes.' (Hutto & Myin, 2013, 13-14)

'REC assumes that some cognitive attitudes are contentful in the restrictive sense of possessing correctness conditions. REC holds, for example, that we sometimes think thoughts that refer to things beyond themselves – thoughts that can be true or false. Nevertheless, it denies that having thoughts with content – so understood – is fundamental to all cognition.' (Hutto & Myin, 2017, 12)

Now, one way in which 'Radical' E-theorists can support their preferred non-representational approach is simply to illustrate how various E-theoretic tools are able to explain various intelligent/cognitive abilities without appealing to representational mental mechanisms or ascribing representational content. These kinds of illustrative case studies, if successful, would show that the notion of internal/sub-personal representation is *not compulsory* when providing explanations in cognitive science. Some of the best-known examples of such *non-representational/non-computational* theories in the literature include: Rodney Brooks' (1991) work in robotics using 'subsumption architecture' instead of internal symbolic processing; Van Gelder's (1995) explanation/description of James Watt's 'Centrifugal Governor' as a dynamic system whose cyclical, self-organizing behaviour is governed by differential equations rather than any kind of sequential, algorithmic process; work by Kelso (1984, 1995) and Beer (1995, 2003) showing how some (simple) human behaviour can likewise be modelled as dynamic systems; work by Turvey et al (1981) which develops Gibson's (1979) notions of 'ecological optics' and 'optical flow'; O'Regan & Noe's (2001) sensorimotor account of visual perception; Freeman & Skarda's (1985, 1990) work on olfaction⁴.

An alternative strategy for Radical E-theorists would be to go on the attack and provide philosophical/theoretical arguments *against* the 'conservative', representational approach.

Anthony Chemero (2009) explicitly advises that we should prefer the former strategy, devoting the opening chapter of his book to warning us away from the temptation of making a priori philosophical arguments with empirical conclusions about how the mind must work:

³ Garzón (2008) likewise uses the term 'radical' to characterize his own anti-representational theory.

⁴ Note: some of these authors do not completely eschew all talk of sub-personal representations – e.g. O'Regan & Noe. But they do all propose explanations of mental phenomena that do not substantively rely on internal representations or computations defined over them. Thanks to an anonymous referee for this journal for pressing me to clarify this point.

‘Radical embodied cognition, claim 1: Representational and computational views of embodied cognition are mistaken.

Radical embodied cognition, claim 2: Embodied cognition is to be explained via a particular set of tools, T, which... do not posit mental representations.

Although I think that claim 1 is correct, I will not argue for it. The point of chapter 1 is that no one believes arguments like that.’ (Chemero, 2009, 29)

Having stated this methodological stricture, Chemero pursues the former argumentative strategy by illustrating how various mental phenomena and abilities can be explained without recourse to representational content. Moreover, he predicts that ‘experimentally discovered facts will largely determine the appropriate theoretical approach in cognitive science’ (ibid, 16). And thus, despite proudly flying the ‘radical’ flag, Chemero ultimately endorses a kind of ‘theoretical pluralism’ (ibid), advising that ‘we should allow many theoretical flowers to bloom’ (ibid) – which, presumably, includes allowing for the possibility of representational flora.

In contrast, Dan Hutto & Erik Myin (2013, 2017) in their pair of books defending a Radical E-theoretic approach (REC), pursue the latter strategy, offering various philosophical arguments for a stronger kind of anti-representational claim: that the very idea of ascribing representational contents to neural or sub-personal mechanisms is a *theoretical confusion*. This whole content-involving approach to cognition (CIC) is, they allege, riddled with ‘intractable theoretical difficulties’ (Hutto & Myin, 2017, 39) – the most fundamental of which is supposed to be the impossibility of solving the Hard Problem of Content (HPC).

‘...defenders of CIC face an intractable problem, which we will call the Hard Problem of Content. This problem arises when they try to make their favoured notion of informational content credible to explanatory naturalists.’ (Hutto & Myin, 2013, 21)

And again in their more recent book, this Hard Problem of Content is glossed as:

‘...explaining how it is possible to get from informational foundations that are noncontentful to a theory of mental content using only the resources of a respectable explanatory naturalism.’ (Hutto & Myin, 2017, 41)

The stronger anti-representational conclusion then is that the notion of representation can never, at least at the neural or sup-personal level, provide any useful explanations in cognitive science. This stronger claim is perhaps somewhat surprising given that Hutto & Myin, like Chemero before them, repeatedly advise against a priori, ‘armchair’ philosophical theorizing, and also given that they claim that the disagreement between REC and CIC is ultimately supposed to be settled in the ‘empirical court’ (Hutto & Myin, 2017, 19).

I am not sure that I share these sorts of general concerns about the legitimacy of a priori ‘armchair’ theorising⁵. Indeed, it is not clear to me why ‘armchair’ theoretical reflections should not count as one more species of flower that we ought, following Chemero’s advice, allow to bloom⁶. Nevertheless, the point of this paper will be to criticise Hutto & Myin’s

⁵ As an example that is relevant to the present topic: William Ramsey’s (2007) important book-length critique of representational assumptions in cognitive science is largely concerned with conceptual or theoretical worries about the very idea of a mental representation (an internal representational vehicle), only briefly mentioning positive/concrete non-representational explanations, such as ‘Dynamic Systems Theory’, towards the end of the book.

⁶ The phrase ‘Let a thousand flowers bloom’ is due originally to Mao Zedong, who used the verse ‘Let a hundred flowers bloom; let a hundred schools of thought contend’, in a number of speeches in 1956. It is a matter of scholarly dispute whether this superficially liberal policy was intended from the start as a means of flushing out dissidents and critics of communist party. In any case, given the subsequent

philosophical arguments for their stronger anti-representational conclusion. For what its worth, I am sympathetic to the *weaker* kind of anti-representational claim – that we ought not assume that all psychological theories/explanations have to be given in terms of states or mechanisms with representational content and that at least some mental phenomena or abilities are in fact better explained in non-representational terms. I am also pretty sympathetic to the sort of broadly Wittgensteinian approach that Hutto & Myin gesture towards in answer to the ‘Hard Problem of Content’ (see section 2, below). And so I hope that these critical remarks will be taken in the spirit that they are intended – that is, from a sympathetic ‘fellow traveller’ who shares some of Hutto & Myin’s concerns about the over-proliferation of representational ascriptions.

The plan for the rest of this paper then is as follows: in section (2) I briefly discuss Hutto & Myin’s own socio-pragmatic answer to the HPC; in section (3) I discuss their accusation that CIC suffers from ‘intractable problems’ dealing with mental causation; in section (4) I turn to Hutto & Myin’s closely related criticisms of teleo-semantic theories of content; in section (5) I consider the theoretical problems that they try to raise against representational accounts of imagination. In each case I will suggest that Hutto & Myin’s arguments for the stronger anti-representational position are not persuasive and that we should therefore prefer the weaker anti-representational stance. In section (6) I suggest that in a number of cases the disagreement between Hutto & Myin and their representational opponents threatens to become a merely verbal dispute and that this worry provides yet another reason to prefer the weaker to the stronger anti-representational thesis. Section (7) provides a short summary and conclusion.

(2) HPC: Dissolution or Solution?

It can be somewhat unclear whether Hutto & Myin take themselves, in championing REC, to be offering a straightforward *answer* to the hard problem of content or whether they take themselves to have *no need to answer* such a problem, or perhaps pseudo-problem, in the first place. In other words, it can be unclear whether they take themselves to be solving or dissolving the hard problem of content. For example – at various points Hutto & Myin use phrases like ‘REC avoids’ or ‘walks away from’ having to deal with the hard problem of content. And of course, to the extent that one is offering a theory or explanation of some specific phenomenon that makes no mention whatever of representational content, then clearly one will not have to deal with any troubling philosophical worries about the natural grounds or determinacy of content.

However, elsewhere in the book Hutto & Myin explicitly allow that:

‘all explanatory naturalists ... must ultimately face up to the HPC one way or another’ (p41).

Since, to repeat, REC does *not* deny the existence of representational mental states altogether, it too must ultimately answer the question: why/how/in virtue of what do these representational states get their determinate representational content?

There are two parts to the story that Hutto & Myin want to give, on behalf of REC, in answer to this question. On the one hand they appeal to a basic kind of non-representational but still ‘intentional’ property – which they call “teleo-semiotic intentionality” – that can be possessed

history of ‘Anti-Rightist’ purges under Chairman Mao’s regime, the phrase now often carries a skeptical or ironic edge – though whether Chemero intended such an ironic reading I cannot say!

by simple organisms with ‘basic minds’⁷ and also by sub-personal neural mechanisms. (We’ll come back to teleo-semiotics in section 4, below.) On the other hand we are told that ‘[full-blown, truth-conditional] content only arises when special sorts of sociocultural norms are in place... involving the use of public symbol systems’ (145) – and we are directed to a paper by Hutto & Satne 2015 for further details. This paper, as Hutto & Satne explicitly allow, only really gestures towards a possible future account of the origins of content as arising from social tendencies to conform and copy and to teach and enforce adherence to norms – an account that would be in the general spirit of Wittgenstein, Brandom and the later Davidson with its appeal to social practices and to something like ‘triangulation’ between subjects as essential to content and objectivity. This is not the occasion to try to flesh out or scrutinise the details of such a social/pragmatist answer to the ‘Hard Problem of Content’ – indeed, I think it’s fair to say that there aren’t really many details given either in Hutto & Myin or in Hutto & Satne. But the general picture is what you might call a ‘top-down’ account insofar as the primary locus of content is taken to be a whole person, embedded within a wider normative, rule-governed practice in a community. For example, Hutto & Myin approvingly quote⁸ Rosa Cao’s verdict that ‘the real epitome of a semantic information using system is not the brain or parts of it, but the whole organism’ (Cao, 2012, 70).

Now, as I say, I am actually pretty sympathetic to this general kind of picture. But for present purposes my aim is neither to criticise nor to defend it. The point is just that *supposing* we were to accept something like this social/pragmatic sort of answer to the HPC – i.e. to the fundamental source/origin of content – it is not clear what motivation there would then be to try to definitively rule out ascribing representational content to some sub-personal structures/processes. For it seems that it is at least possible that it could be useful or explanatory when investigating how these internal structures/processes combine to produce intelligent, norm governed behaviour, to attribute them *derivative* representational contents – just as we attribute representational content to written sentences, to public signs and symbols, to recorded public announcements etc., when explaining our social, linguistic behaviour. Of course, in a sense, the content of these internal structures/processes would be derivative or dependent on the more fundamental social/normative practices, whatever those are exactly. But that need not make their representational content any less determinate or genuine or fit to figure in explanations than the content of printed sentences or public signs.

Something like this sort of ‘top-down’ picture is, I take it, the approach that Dennett (1987, 1991, 1995) has famously championed. Whilst the fundamental locus of intentionality for Dennett is a whole subject, he thinks it can still be useful to ascribe content to various sub-personal mechanisms – or ‘homunculi’ – whose behaviour can also be interpreted from the ‘intentional stance’ and whose functioning together combines into the production of the whole subject’s representational mental abilities.

To take a favourite example of Dennett’s⁹, consider an electronic calculator: let’s just assume that the numerals and symbols that the device displays only have determinate meaning insofar as it is used as part of the human practice of mathematics and counting etc. In other words, we’re assuming that the ultimate source of intentionality here is ‘top-down’ and derived from human/social norms. (So if a physical duplicate of the device appeared out of nothing by sheer random quantum chance somewhere in outer space, it would *not* count as calculating or representing anything.) But still, when trying to figure out how the electronic calculator works or give an account of how all its internal parts function so as to produce the numerals

⁷ See Korbak (2015) for an argument that a consequence of Hutto & Myin’s way of distinguishing between basic and non-basic minds is that there would be no basic minds, for even simple biological signaling systems would count as using language/symbols.

⁸ See Hutto & Myin, 2017, 238.

⁹ The first occurrence of this example that I can find in Dennett’s oeuvre is Dennett 1982, 224. See also: Dennett 1987, 78, Dennett 1998 315-316, Dennett 2013, 96-97.

and symbols that it outputs, it can be overwhelmingly useful to speak of internal *parts* of the calculator in the following way: this is the bit that does the subtracting, this is the bit that carries the one, this is the bit that ‘remembers’ the result of previous steps in the calculations etc. Eventually, when we, as it were, zoom in far enough, there will presumably come a point when it ceases to be useful to think in representational terms and we will describe these microscopic parts in purely physical/mechanistic/functional terms. But that does not obviously show that representational ascriptions at ‘higher levels’/scales were not genuine or not explanatory or were a mere *façon de parler*. (We will return to this point in section 3, below.)

Of course, by now we are all very familiar with the various ‘non-classical’ approaches in cognitive science – from connectionist architecture to dynamical systems theory – which show that we ought not *assume* that there *must be* neatly decomposable mechanisms in our brains, as there are in an electronic calculator, which can be attributed contents in this way. But nor is there any obvious motivation to rule out in advance, on philosophical grounds, that this could ever turn out to be the case. Let’s wait and see! In other words, even supposing that we grant Hutto & Myin their preferred social/pragmatic solution to the ‘Hard Problem of Content’, there would be no obvious reason to think that attributing *derivative* contents to neural/sub-personal mechanisms would have to be a ‘theoretical confusion’. Hutto & Myin’s own positive picture then so far provides no reason to go beyond the weaker anti-representational position to the stronger one.

(3) Mental Causation

A key complaint that Hutto & Myin make against content-involving approaches (CIC and CEC) is that they suffer from ‘intractable problems’ dealing with mental causation, because the representational contents posited by these theories are bound to be unable to provide the sort of *causal* explanations sought by cognitive science.

“...we are owed an account of how the abstract content of a concrete vehicle could make any causal difference to cognition.” (Hutto & Myin, 2017, 35)

Hutto & Myin repeatedly claim that such causal explanations will already be provided by non-representational features – e.g. by appealing to a ‘functional isomorphism’ we can already explain how some neural structure/process allows for useful guidance of some kind of behaviour without mentioning content – and so the supposed representational content would be ‘screened off’ from playing an explanatory role. The positing of representational content would be redundant.

Let me now make two critical points about the general form of this complaint.

(i) There is a very extensive philosophical literature about causation, reduction and the relations between ‘different levels’ of explanation¹⁰. One key issue, very briefly, is that from the point of view of fundamental physics, it can seem that no other non-fundamental properties are needed nor could they do any real causal work/explanation. And yet we go about making apparently true and useful causal claims couched in terms of higher-level properties (chemistry, biology, economics etc.) Hutto & Myin make clear that they have no sympathy with the sort of extreme reductionism that treats fundamental physics as the only legitimate/genuine level for naturalistic, causal explanation¹¹. But Hutto & Myin’s complaint about causal inefficacy – that representational properties are redundant, having been screened

¹⁰ See e.g. Nagel (1935, 1951), Fodor (1974), Putnam (1975), Cartwright (1983, 1999), Kim (1990, 1992), Dupré (1993), Rosenberg (2006).

¹¹ See Chapter 6 of Hutto & Myin 2017, especially p124-126.

off by other non-representational properties – leaves it unclear why *representational* properties *in particular* should be considered more problematic than many other non-fundamental but scientifically respectable properties. After all, the fact that a physical system can, in theory, be described and its behaviour explained and predicted using only, say, the language of physics and chemistry, without mentioning biological features like species or ecological niche or chromosome etc., does not – unless you are an extreme sort of reductionist – show that biological level explanations which mention such features are not also perfectly kosher causal claims. And so even if we grant that there is some non-representational level of description, which describes the causal role of various physical structures (perhaps both internal and external) in producing an animal's behaviour without mentioning content – e.g. mentioning only, say, functional isomorphisms or dynamic embodied coupling – this does not obviously show that there is not some other level of description at which we can *also* make useful explanatory claims which *do* mention content.

Think here again of the electronic calculator – sure, we can (in theory) describe it at the level of atoms and molecules and thus make successful predictions and explanations of its behaviour without mentioning representational content or addition or subtraction etc. But unless one is already committed to some large metaphysical thesis that only the most microscopic or fundamental level is real or can feature genuine causation etc., then this is so far no reason to rule out that there might *also* be useful descriptions, explanations and *ceteris paribus* predictions which describe the sub-components of the calculator in representational terms.

(ii) Moreover, even supposing that there *is* a *special* problem with appealing to representational properties in causal explanations, a problem that is *not* faced by, say, biological or economic or geological properties, it is hardly a problem that REC avoids. To repeat: Hutto & Myin stress that REC allows that there are indeed at least some contentful mental states. And so their overall theoretical picture will still have to deal with this same range of familiar issues concerning mental causation that any other non-eliminativist picture faces – viz. it intuitively *seems* that the representational contents of my thoughts and desires etc. are relevant to their causing and producing the specific behaviour they do and yet it is hard to understand how such representational properties really can be doing causal work. Gesturing towards the idea – and to repeat: I think it is a potentially *good* idea – that content essentially depends on the subject's being embedded in social-normative practices does not obviously offer any kind of solution to this familiar kind of worry about mental causation.

In summary then, unless some rather more heavy-duty metaphysical theses about causation and reduction are added into the mix, Hutto & Myin have not shown that content-involving approaches suffer from 'intractable problems' dealing with mental causation. And even if there are intractable problems in the vicinity here, they would seem equally to be problems for Hutto & Myin's own preferred view. And so again it seems we do not yet have a reason to prefer the stronger brand of anti-representationalism to the weaker.

(4) Teleo-Semiotics vs. Teleo-Semantics

A more specific version of Hutto & Myin's complaint that ascribing representational content to internal structures cannot provide for *causal explanation* is levelled at Teleo-Semantic accounts of such content¹²:

“...teleological theories understand the intentional by appeal to proper functions, and when it comes to understanding proper functions, history sets the standard and supplies the norms. It is history that determines whether a current attitude represents correctly or incorrectly or what

¹² See e.g. Dretske (1981, 1995), Millikan (1984), Papineau (1984). Neander (1995)

it is directed at. Consequently, understood via teleological theories, intentional properties are debarred from playing any kind of causal role in the synchronic production of intelligent behaviour

...The intentional, if understood in terms of biological functions, does no causal, mechanistic work.” (108-9)

And in a footnote¹³ Hutto & Myin commend Ruth Millikan for explicitly stating: ‘that a thing has a teleofunction is a causally impotent fact about it’ (Millikan 1993, 186).

But Hutto & Myin don’t wish to abandon teleological explanations of intentionality altogether. Quoting Peter Godfrey-Smith (2006, 60) they write:

“all is not lost – appealing to biological functions can still explicate ‘an important kind of natural involvement relation’” (114)

And then a couple of pages later:

“REC tweaks the standard teleo-semantic story by dropping any commitment to understanding these sorts of basic intentional relations as involving semantic properties such as reference or truth.” (116)

Which brings us onto the other part of Hutto & Myin’s own preferred story about the origin/source of mental content – their notion of primitive/basic teleo-*semiotic* intentionality.

Taking their inspiration from Brentano’s (1874 [2009]) seminal remarks that intentional states are those which are directed at a target object, Hutto & Myin insist that the most basic forms of intentionality – ‘Ur-intentionality’ – are states or behaviours which somehow point beyond themselves towards some targeted object, but which do not have anything like correctness conditions or truth-conditions and are therefore intentional but not representational (as Hutto & Myin are using these labels). Of course notions like ‘pointing at’, ‘directed at’, ‘targetting’ etc. are really just metaphors that leave all the hard work still to be done. A flag in high wind points in the direction of the prevailing wind, a shaft of sunlight penetrating through the cloud is directed at/onto some object – but these are not supposed to be counted as intentional phenomena. So what does distinguish intentional from non-intentional? According to both Hutto & Myin’s teleo-semiotics and the more orthodox teleo-semantic view, we do this by appealing to biological function. Mental processes or inner states have genuine, original intentionality as they have the biological function of tracking a specific kind of distal, environmental feature. According to teleosemantic theories, the inner state/mechanism S represents that F is present because tracking instances of F contributed to the fitness of the animal’s ancestors and thus explains why S continues to exist in the present animal. Likewise on Hutto & Myin’s teleo-semiotic view, the inner state/mechanism is *targeted* or *directed at* F’s because such targeting contributed to fitness and so is the reason why the state/mechanism persists in the species to this day.

According to teleo-*semantics*, if S is activated in the absence of F, the state is then *misrepresenting* how things are out in the environment. Now of course Hutto & Myin want teleo-*semiotics* to eschew all talk of correct and incorrect representing, but they still need biological proper function to draw a line between all the possible circumstances where the state/mechanism is directed onto its proper biological *target* and all those other circumstances where it is directed onto something else. For Hutto & Myin do not want the teleo-semiotic object/target of a frog’s tongue-snapping mechanism¹⁴ to end up being a highly disjunctive

¹³ See fn. 15 in Chapter 5 of Hutto & Myin, 2017.

¹⁴ The well-worn example of a frog snapping its tongue at flies derives originally from the classic paper ‘What the Frog’s Eye Tells the Frog’s Brain’ by Lettvin, Maturana, McCulloch & Pitts (1959).

kind: ‘flies *or* shadows *or* BB pellets *or* etc.’ So by appealing to the fact that it is successfully targeting flies which explains the tongue-snapping’s contribution to fitness, whereas the occasions where the tongue-snapping was directed at mere shadows etc. did not contribute to fitness, teleo-semiotics, like teleo-semantics, seeks to avoid problems with indeterminate or highly disjunctive classes.

So whilst Hutto & Myin insist that teleo-semiotics draws a line between *aligned* and *misaligned*, or, *appropriate* and *inappropriate* responses – rather than between correct and incorrect, or, accurate and inaccurate responses – the appeal to biological function serves an entirely similar purpose as in teleo-semantics:

“How can we understand the sort of world-directedness exhibited by contentless Ur-intentionality in naturalistic terms? The short answer is that this can be done by RECTifying teleosemantics – by stripping it of its problematic semantic ambitions and putting its basic apparatus to new and different theoretical use.” (104)

‘Why... should we think the frog is targeting flies and not all the other items on the long list of things it is disposed to chase? ...the biological function of tongue-snapping behavior is not disjunctive at all; it is directed at flies and flies alone. This is because ingesting flies, as a matter of fact, met the needs of this type of frog’s ancestors.’ (115)

Let me now make three points about Hutto & Myin’s proposed teleo-semiotics:

(i) Firstly: insofar as it is true that teleo-semantic content would be causally impotent – as both Hutto & Myin and Millikan claim – then surely teleo-semiotic intentionality must likewise suffer from causal impotence. Hutto & Myin’s earlier complaint that “The intentional, if understood in terms of biological functions, does no causal, mechanistic work” would equally apply to their own account of Ur-intentionality since they explicitly appeal to *historical* biological functions to determine the intentional target in exactly the same way that teleo-semantics determines representational content.

If causal impotence is a problem for Millikan, Papineau and Dretske, why is it not likewise problematic for Hutto & Myin? Or vice-versa: if appealing to historical biological facts is not a problem for teleo-semiotics, surely it should be likewise unproblematic for teleo-semantics.

(ii) Secondly, appealing to biological proper function is meant to partition the space of all possible occasions where the state/mechanism is activated into those that are ‘appropriately directed at the intentional target’ and those that are not – just as in teleo-semantics. So appealing to evolutionary function may allow us a principled basis to rule out such possible triggers of the frog’s tongue-snapping as: shadows and BB pellets and poisonous flying insects that look like flies etc. from forming part of the intentional target of the mechanism/behaviour.

But even after we have ruled out these triggering circumstances from determining the intentionality of the state/mechanism, familiar worries about indeterminacy remain that have been extensively discussed in the literature on teleo-semantics¹⁵. So even if we can rule out shadows and BB pellets etc. from forming part of the content, or being part of the intentional target, appealing to biological function does *not* apparently allow us to distinguish between ‘flies’, vs. ‘nutritious flying insect’ vs. ‘flying food’, vs. ‘visible flying food’, vs. ‘visible, non-poisonous flying food’, vs. ‘ingestible item that can be caught via the tongue and which enhances fitness’ etc. The standard worry here is that the intentional content/target of *all*

¹⁵ See e.g. Fodor 1996, Fodor & Piatelli-Palmarini 2010, Neander 1995, Price 1998, Millikan 1991, Papineau 1997.

evolved mechanisms will end up simply being: *fitness enhancer* as any other property or feature is bound to be strictly less relevant to reproductive fitness.

At times Hutto & Myin seem to imply that such indeterminacy doesn't matter when it is a primitive intentional target/object rather than a representational content¹⁶. Likewise, Hutto & Satne write:

“Understanding the range of things to which organisms respond as a disjunction of counterfactual items is a feature of Ur-intentionality, not a problem for it. Not being a theory of intentional content familiar worries about the determinacy and specificity of content do not apply.” (Hutto & Satne, 2015, 531)

But why should it be any less worrying that the intentional *target* of the frog's tongue snapping mechanism is indeterminate/disjunctive than that the representational *content* of its snapping mechanism is indeterminate/disjunctive? To repeat the issue: so far as evolutionary theory is concerned, we could equally well say that ingesting edible flies or ingesting flying food or ingesting non-poisonous flies or ingesting edible fitness enhancers, etc. was what met the froggy ancestral needs. Insofar as this is a genuine worry, teleo-semiotics will presumably be unable to distinguish between any of these as the 'appropriate' intentional target of the frog's tongue-snapping mechanism. Notice also that these are not all co-extensive categories – the set of flies is not identical to the set of edible, non-poisonous flies – so it is not just a matter of the tongue-snapping being directed at flies 'under any description' or under any mode of presentation, as Hutto & Satne (2015, 530-531) claim.

Now perhaps Hutto & Myin would hold that they are only giving an account of a very basic/primitive form of intentionality, for which this sort of indeterminacy is not such an issue – and then point out that more determinate kinds of intentionality are given a different explanation in their story, by appeal to socio-cultural normative practices etc.

But then this same move should also then be legitimate for a defender of teleo-semantic content: oh sure, at the most basic, primitive level, the content is indeterminate between representing that there is a fly, or flying food, or a flying fitness enhancer etc. But once the system becomes more complex and gets something like recursive conceptual structure and full blown beliefs and intentions, well then the content can become much more determinate and fine-grained. And indeed this is, I take it, just the sort of move that Papineau (1997, 2016) makes in response to these sorts of indeterminacy worries for teleo-semantics.

(iii) Thirdly: another famous, oft-discussed objection to teleo-semantics is, of course, Davidson's (1987) Swampman case. We are to imagine that an atom-for-atom duplicate of a normal human has arisen by sheer chance – lightning hitting a swamp say. According to teleological theories, such a duplicate could not possess any kind of contentful mental states – it could not have any meaningful thoughts or desires or intentions – *despite* being physically identical and therefore behaving and talking identically, because this duplicate would have no evolutionary history and so no proper biological functions. But surely, so the intuition is supposed to go, such a creature *would* have thoughts and desires etc.

Clearly this Swampman worry – insofar as it is a genuine worry – would apply just as much to teleo-semiotics. Swampman could not have any mental states or inner mechanisms with intentional *targets* either. And nor could swamp honey bees or swamp magnetosome bacteria etc.

In the same vein, consider another classic problem case, due to Paul Pietroski (1992), of the snorfs and the kimus:

¹⁶ See Hutto & Myin, 2017, p115-116.

“The kimus live near a large rocky hill. Their only predators are snorfs, carnivores who roam past the hill each morning. Kimus used to be ‘color-blind’. But in virtue of a genetic mutation, one particular kimu – call him Jack – came to have an internal mechanism M that produced tokens of a physically specifiable state type B in the presence of certain wavelengths of light. Each morning, something red on the hilltop caused Jack to form a B-token when he looked up. And Jack (like his descendants) turned out to have a ‘fondness’ for red things; i.e. other things being equal, Jack would move towards the distal causes of B-tokens when such tokens were produced. So each morning, Jack trudged up the hill and thereby avoided the snorfs. Natural selection took over; and Jack’s mechanism type proliferated throughout the species...

There are several candidates for the content of a (current) kimu’s B-token: Lo, redness; Lo Wavelength W; Something nice is over there; There’s that nice mountaintop again; etc. Or perhaps there is no determinate content of B-tokens to speak of. But one thing strikes me as obvious: B-tokens are not about snorfs. Nonetheless this is the consequence that Milikans’s account delivers.” (Pietroski, 1992, 273)

After all, by hypothesis the kimu’s visual system is sensitive to the *redness* of the dawn sunlight. Indeed after the adaption has spread through the population it might be that no kimu has actually seen or encountered a snorf for tens of thousands of years – the snorf might long ago have become extinct! But notice again that Hutto & Myin’s teleo-semiotic account will likewise deliver the counter-intuitive verdict that the intentional *target* of the kimu’s B-token is something like: *snorf-free area*, or *predator-free place*.

Now, I do not wish to suggest that these much-discussed objections are definitely fatal to teleo-semantics or to teleo-semiotics¹⁷. The point is just that Hutto & Myin repeatedly emphasise – quite legitimately – how teleo-semantics faces various theoretical problems when trying to ground content in biological function. But these same problems are presumably bound to be inherited by teleo-semiotics insofar as it too seeks to ground intentionality in biological function. Given that teleo-semiotics is just a ‘tweak’ on teleo-semantics, it should be unsurprising that it will face much the same problems about causal impotency, indeterminacy, swampman etc. And to the extent that teleo-semiotics can escape these worries, then it is unclear why teleo-semantics cannot also make a parallel escape.

Hutto & Myin’s discussion of teleological theories of intentionality does not provide a compelling reason to take the stronger anti-representational approach – at least as such an approach is developed by Hutto & Myin. For the theoretical problems they raise for a teleological theory of content would seem to apply equally to their own non-representational teleological theory of intentionality.

(5) Imagination, Content & ‘Direction Of Fit’

A more specific anti-representational complaint that Hutto & Myin (2017) make concerns the ascription of representational content to *imaginative* mental states. Though this is not such a central plank in their anti-representational case, I think it reveals an important distinction concerning the notion of correctness or satisfaction conditions that I think they have overlooked or mishandled.

Drawing on work by Langland-Hassen (2015), Hutto & Myin suggest that it is hard to see how imaginative mental states could have correctness/accuracy conditions. They claim that the two main ways of trying to ascribe correctness conditions to a sensory imaginative state is to treat it either as: representing that the imagined state of affairs actually, presently obtains – which they call representing ‘as present’, or, (ii) representing that the imagined state of affairs

¹⁷ See Papineau (2001) for a response to the Swampman worry on behalf of teleosemantics.

is merely possible on some or other notion of modality – physical, nomological, metaphysical, etc – i.e. which would be representing ‘as possible’.

The problem then is supposed to be that on the representing as present view, imaginative states are almost always misrepresentations and it is supposedly implausible that there could be ‘a useful cognitive faculty that continually issues in misrepresentations’ (Langland-Hassen, 2015, 671). On the other hand, on the representing as possible view, imaginative states will be almost always be trivially successful and this is supposedly a problem as then we would no longer be able to capture the distinction between useful and non-useful imaginings so far as action guidance and planning are concerned.

Now, no doubt there are various ways that this Scylla and Charybdis of ‘too little accuracy’ or ‘too much accuracy’ might be resisted by someone who wanted to defend the idea that imaginative state have correctness conditions¹⁸. But for present purposes let’s just grant, for the sake of argument, that Langland-Hassen’s criticisms are correct. The conclusion that Langland-Hassen draws is that we should not think of imaginative mental states as having either ‘as-present’ or ‘as-possible’ correctness conditions¹⁹.

But Hutto & Myin suggest an importantly *different* conclusion:

“In light of the failure of both the as-present and as-possible views, one option is to give up on the idea that imaginings represent anything at all.” (187)

And indeed for the class of ‘basic sensory imaginings’²⁰ this is the option that Hutto & Myin endorse. But it is not at all obvious why we should move from the claim that imaginings lack correctness conditions to the conclusion that they are entirely non-representational.

There are a host of propositional attitudes which do not have correctness conditions but which are straightforwardly representational – e.g. desiring that p, fearing that p, intending that p etc. Some of these attitudes are often said, in Anscombe’s (1957) phrase, to have the opposite ‘direction of fit’ to beliefs and judgements. E.g. a desire or an intention aims to change the world so as to fit the content of the mental state, whereas we form and change our beliefs and judgements so that their contents match the world. States like desire or intention then are said to be satisfied or unsatisfied rather than correct/incorrect or accurate/inaccurate.

But it is also very plausible that some propositional attitudes have *neither* direction of fit. For example: interrogative/questioning attitudes such as: wondering whether p; or consider the attitude of suspending judgement whether p. These attitudes are straightforwardly

¹⁸ An anonymous referee for this journal points out that the notion of an *efference copy*, which has been employed in psychological theories of motor control going back at least to Von Helmholtz, seems to provide a clear counter-example to the claim that it is implausible that a useful cognitive faculty could mostly traffic in misrepresentations. For in some circumstances efference copies will mostly misrepresent the movement-producing signal but this misrepresentation is precisely how the system manages to control the movement. See e.g. Wiese (2017) for discussion of such systematic misrepresentation.

¹⁹ Langland-Hassen ultimately argues that complex, hybrid imaginative states, combining a sensory image component with a propositional attitude component, *do* have correctness conditions. Langland-Hassen explicitly allows that *some* imaginative states may be propositional despite lacking any direction of fit – ‘just as there are cases where a person entertains a proposition while remaining neutral about its truth’ – though he rejects the conclusion that the class of imaginative states *in general* lack satisfaction conditions. See Langland-Hassen, 2015, 673-4.

²⁰ Hutto & Myin, like Langland-Hassen, distinguish between ‘basic sensory imaginings’, which they take to be non-contentful, and more complex, ‘hybrid’ imaginative states, which are partially constituted by propositional attitudes, such as intentions, and which they therefore allow can possess content. See Hutto & Myin, 2017, 188-189.

representational – they have a propositional content that *p*. But they do not seem to have either direction of fit – their propositional content is neither supposed to match the external world, nor to be a blueprint for changing the world or bringing about the state of affairs that is represented. Likewise then it seems entirely possible to hold that an imaginative state can have a representational content that *p*, without the state thereby having either correctness conditions or satisfaction conditions – for imaginative states are plausibly just one more type of propositional attitude that have content but have no direction of fit.

Hutto & Myin write:

“The claim that the imagination is contentless across the board has some prominent supporters. Searle (1983), for example, argues that imaginings, as a class, lack a direction of fit. Accordingly, imaginings must be silent about how things stand with the world; they are – to borrow McGinn’s (2004, 21) phrase – neutral about reality.” (188)

Lacking a direction of fit and lacking representational content should not be conflated like this. Whether a mental *state* can be evaluated as correct or incorrect is one question – whether it has a propositional content is another. My *mental states* of desiring that *p*, or wondering whether *p*, or suspending judgement whether *p*, all lack correctness conditions – but we can still perfectly well ask what state of affairs they all represent in common and whether that state of affairs or proposition, *p*, in fact obtains or not.

At the risk of belabouring the point, let me briefly run through an analogy. Consider a drawing or picture of a house. This physical picture might be used on Monday as a map or depiction of some particular house – and thus the drawing can be evaluated as accurate or inaccurate in various ways according to the state of the actual house. On Tuesday I might use the very same drawing as a blueprint to build a new house – and so now it can be evaluated as satisfied or not depending on how well the builders manage to meet the specifications of the drawing when building the new house. On Wednesday I might use the drawing to pose a question – is there a house in Bochum that is like this? Who would live in a house like this? Now the drawing is being used neither to represent how things are nor to represent how I desire them to be. And on Thursday perhaps I use the drawing as an object of pure aesthetic contemplation. But on all four days, in all four roles, the drawing has a common representational content throughout – it represents the same possible type of building²¹.

Hutto & Myin’s discussion of imagination and correctness conditions then does not reveal an intractable theoretical problem with attributing content either. Whatever theoretical problems may exist with the notion of ascribing correctness conditions to imaginative states, this is not yet any obvious reason to conclude that these states lack representational content altogether

(6) Merely Verbal Disagreement?

The final issue I want to consider is one that I have heard voiced a number of times in conversation, and one which I know that Hutto & Myin are well aware of²². This is the

²¹ Which is *not* to say that it explicitly represents *that* the building is possible.

²² The only instance I have been able to find of this worry being lodged against Hutto & Myin in print is Rolla (2017, 52). The same worry about merely verbal disagreement is lodged against Chemero in Reboul (2011). There is an interesting online exchange between Hutto and Gualtieri Piccinini on ‘The Brains Blog’, where Piccinini voices the worry about merely verbal disagreement:

<http://philosophyofbrains.com/2014/12/11/radicalizing-enactivism.aspx>

Hutto & Myin intimate that they are well aware of the worry in the preface to their second book, where they write: ‘The changes REC aims to install in the way we think about thinking require theoretical adjustments to our conception of cognition, not mere verbal tweaks.’ (H&M, 2017, xii-xiii).

suspicion that at various points Hutto & Myin may be in danger of getting into a *merely verbal* dispute.

At a number of points, Hutto & Myin claim that some avowedly representational psychological or philosophical theory is *almost correct except* for its misguided commitment to treating the structures/processes in question as representational. So for example: De Brigard's (2014) account of memory is said to be 'mostly correct, sans its commitment to the ...claims that information is encoded and processed' (28). And concerning Kandel's (2001, 2009) account of memory and learning, Hutto & Myin comment: "REC does not question the details of Kandel's explanations of how experience leaves a trail of structural changes in the brain... it challenges the standard cognitivist interpretation" (239). Similarly, when discussing Predictive Processing (PPC), they claim: 'Reducing uncertainty and reducing free energy can be easily cashed out in REC terms...we can keep what is best in PPC whilst shelving the philosophically confounding talk of [representational content].' (70-71) And we have already seen how Hutto & Myin describe their own non-representational teleo-semiotic theory as a tweak on representational teleo-semantic theories (see section 4, above).

But if so much is conceded to be correct in these theories, the worry arises whether any remaining disagreement is *merely verbal*; that is, a disagreement over whether or not to use representational-*sounding* labels for various elements in an otherwise *agreed* upon theory. We certainly seem here to be a long way from the claims at the start of the book that REC is ultimately to be vindicated in the 'empirical court'.

I think that this worry about merely verbal disputes became especially acute towards the end of their recent book, when Hutto & Myin return to discussing the notion of information, insisting that "information is not *really* picked up, or passed on, or pooled in the brain for use by the brain" (238). But then in the very next sentence they state that:

"Neurodynamics, understood a la REC, takes the form of informationally sensitive, well-connected neural activity that plays influencing and mediating, as opposed to representational roles..."

Any difference between the claim that neural activity is 'picking up information' vs. the claim that neural activity is 'informationally sensitive' seems, to my ears anyway, rather worryingly slight and verbal.

Likewise, Hutto & Myin critically discuss the common talk in cognitive science of internal or neural 'maps':

"...we can ask whether the brain literally consults the positioning system like a map, model or guide to the navigational possibilities of the environment. That would be to picture the brain as a homunculus, leading to well-known trouble. It would be better to say that the positioning system has properties that are map-like and those properties are exploited in extremely complex ways..." (240-241)

Hutto & Myin insist that "There is a significant theoretical difference between these two interpretations" (241) – though they also allow that: "The theoretical disagreement between REC and CIC is a refined disagreement" (243). I suspect I will not be alone in worrying that the difference between 'consulting a map' and 'exploiting map-like properties' may turn out to be refined to the point of theoretical insignificance.

Now to be fair to Hutto & Myin, on the very first page of the preface to their book they quote from Richard Feynman's Lectures on Physics – 'Even a very small effect sometimes requires profound changes in our ideas' (Feynman et al. 1963, 1-2) – in order to make the point that 'sometimes revolutions in thought ensue from what may seem, for most practical purposes,

only small or marginal changes to a theoretical framework' (Hutto & Myin 2017, xi). And they return to this point (and to Feynman) again at the end of the book when justifying the theoretical importance they attach to distinguishing between informational sensitivity vs. informational 'pick-up'.

'In conducting some kinds of scientific work, attending to these subtle differences may not matter. But these devilish details surely do matter if we are to get a cleaner, leaner and clearer understanding of the properties of cognition.' (244)

It is certainly true that *sometimes* an apparently slight conceptual/theoretical change can turn out to have very significant empirical consequences. And conversely, it is true that *sometimes* an apparently small experimental discrepancy can turn out to require very drastic theoretical change. These are the events that the history of science mostly focuses on. However, it is surely plausible that many *more* small conceptual/theoretical changes turn out to be of only minor or no theoretical importance. And likewise, presumably, the majority of small experimental discrepancies turn out to be mere noise or to require only minor theoretical adjustment. This larger class of insignificant events in 'normal', non-revolutionary science is, understandably, not what historians of science have tended to dwell on. But they form the normal/default situation in science nonetheless.

And so it seems to me that the onus here remains on Hutto & Myin to provide positive reason to think that these apparently subtle/verbal disputes will, contra appearances, turn out to have significant, far-reaching consequences. To put it another way: Hutto & Myin are effectively making a prediction about the future scientific significance of choosing between: 'informational sensitivity' vs. 'informational pick-up', or, 'mental map' vs. 'map-like property' – and this strikes me as an incautious bet to make.

Now of course if it does turn out that a dispute between Hutto & Myin and their representational opponents is a merely verbal dispute, this result would hardly vindicate those theorists committed to the representational cause either! And thus the spectre of verbal disagreement is, I suggest, yet one more reason to prefer the weaker position that representational forms of explanation are simply non-compulsory to the stronger claim that they are bound to involve theoretical confusion. For in a merely verbal dispute, *neither side* suffers from deep theoretical confusion – the only confusion would be shallow and linguistic.

(7) Conclusion:

I have made a number of critical remarks about some of the stronger anti-representational claims and arguments that Hutto & Myin have advanced. In light of these criticisms, I suggest that they have not yet presented a persuasive case for their claim that content-involving approaches to (basic) cognition are bound to be riddled with 'intractable theoretical confusions'. But I would like to conclude on a more irenic note.

Hutto & Myin employ a whole arsenal of 3-letter acronyms in their books and so I hope I will be forgiven for suggesting that we add one more to that list. Located in-between CEC (Conservative E-accounts), which are committed to giving representational explanations, and REC (Radical E-accounts), which are committed to rejecting such representational explanations, I think that the default position for E-theorists and sympathisers should be: TEC – Tolerant E-accounts of Cognition. TEC agrees with REC that representational content *need not* be appealed to in cognitive-science explanations. But TEC allows, in an open naturalistic spirit, that we may yet discover – perhaps after empirical investigation, perhaps after

theoretical reflection – that *some* neural or sub-personal mechanisms are most helpfully thought of, at least for some explanatory purposes, in representational terms²³.

References

- Anscombe, G. E. M. (1957) *Intention*, Oxford: Blackwell.
- Ashby, R. (1956) *Introduction to Cybernetics*, New York: Wiley.
- Beer, R. (1995) ‘A dynamical systems perspective on agent-environment interactions’, *Artificial Intelligence* 72: 173-215
- Beer, R. (2003) ‘The dynamics of active categorical perception in an evolved model agent’, *Adaptive Behavior*, 11: 209-243
- Brentano, F. (1874 [2009]) *Psychology from an Empirical Standpoint*, Routledge.
- Brooks, R. (1991) ‘Intelligence without Representation’, *Artificial Intelligence* 47: 139-159
- Cao, R. (2012) ‘A teleosemantic approach to information in the brain’, *Biology & Philosophy*, 27(1): 49-71.
- Cartwright, N. (1983) *How the Laws of Physics Lie*, Oxford: Oxford University Press.
- Cartwright, N. (1999) *The Dappled World: A Study of the Boundaries of Science*, Cambridge University Press.
- Chemero, A. (2009) *Radical Embodied Cognitive Science*, MIT Press
- Clark, A. (1997) *Being There*, MIT Press
- Davidson, D. (1987) “Knowing One's Own Mind”, *Proceedings and Addresses of the American Philosophical Association*, 60: 441–58.
- De Brigard, F. (2014) ‘Is memory for remembering? Recollection as a form of episodic hypothetical thinking’, *Synthese* 191(2): 1-31.
- Dennett, D. (1982) ‘Styles of Mental Representation’, *Proceedings of the Aristotelian Society* 83: 213-226.
- Dennett, D. (1987) *The Intentional Stance*, MIT Press
- Dennett, D. (1991) ‘Real Patterns’, *Journal of Philosophy* 88 (1): 27-51.
- Dennett, D. (1995) *Darwin's Dangerous Idea*, New York: Simon & Schuster
- Dennett, D. (1998) *Brainchildren*, MIT Press.
- Dennett, D. (2013) *Intuition Pumps and Other Tools for Thinking*, Norton & Co.
- Dretske, F., (1981) *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
- Dretske, F. (1995) *Naturalizing the Mind*, Cambridge, MA: MIT Press.
- Dupré, J. (1993) *The Disorder of Things. Metaphysical Foundations of the Disunity of Science*, Cambridge, MA: Harvard University Press.
- Feynman, R., Leighton, R. & Sands, M. (1963) *The Feynman Lectures on Physics*, Vol. 1, Addison-Wesley.
- Freeman, W. & Skarda, C. (1985) ‘Spatial EEG patterns, nonlinear dynamics and perception: The neo-Sherringtonian view’, *Brain Research* 10: 147–175.
- Freeman, W. & Skarda, C. (1990) ‘Representations: Who needs them?’, in McGaugh, Wienberger & Lynch (eds.) *Brain Organization and Memory Cells, Systems, & Circuits*, Guilford Press.

²³ Material from this paper was presented at the workshop ‘Evolving Enactivism’ at the Ruhr-University Bochum. I am very grateful to the audience on that occasion for their questions and comments. Many thanks in particular to Dan Hutto and Erik Myin for their generous and helpful feedback. I am also extremely grateful to Peter Brössel, Elmarie Venter and Filippo Vindrola who read an earlier draft and provided many insightful suggestions and criticisms. Finally, I would like to thank two anonymous referees for this journal whose comments substantially improved this paper.

- Fodor, J. (1974) “Special sciences, or the disunity of science as a working hypothesis”, *Synthese*, 28: 77–115.
- Fodor, J. (1975). *The Language of Thought*, Cambridge, Massachusetts: Harvard University Press.
- Fodor, J. (1980). “Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology”, *Behavioral and Brain Sciences* 3: 63-73
- Fodor, J. (1996) “Deconstructing Dennett's Darwin”, *Mind and Language* 11: 246–262.
- Fodor, J. & Piatelli-Palmarini, M. (2010) *What Darwin Got Wrong*, Farrar, Straus & Giroux.
- Garzón, F. C. (2008) “Towards a General Theory of Antirepresentationalism”, *British Journal for the Philosophy of Science* 59 (3): 259-292.
- Gibson, J. J. (1966) *The Senses Considered as Perceptual Systems*, Waveland Press.
- Gibson, J. J. (1979) *The Ecological Approach to Visual Perception*, Boston: Houghton-Mifflin.
- Godfrey-Smith, P. (2006) ‘Mental representation, naturalism, and teleosemantics’, in Macdonald & Papineau (eds.) *Teleosemantics: New Philosophical Essays*, Oxford University Press.
- Hutto, D. & Myin, E. (2013) *Radicalizing Enactivism: Basic Minds without Content*, MIT Press.
- Hutto, D. & Myin, E. (2017) *Evolving Enactivism: Basic Minds Meet Content*, MIT Press.
- Hutto, D. & Satne, G. (2015) ‘The natural origins of content’ *Philosophia* 43(3): 521-536.
- Kandel, E. (2001) ‘The molecular biology of memory storage: A dialogue between genes and synapses’, *Science* 294: 1030-1038
- Kandel, E. (2009) ‘The biology of memory: a forty-year perspective’, *Journal of Neuroscience* 29(41): 12748-12756.
- Kelso, J. (1984) ‘Phase transitions and critical behaviour in human bimanual co-ordination’, *American Journal of Physiology: Regulatory, Integrative and Comparative*, 15: R1000-R1004.
- Kelso, J. (1995) *Dynamic Patterns*, MIT Press.
- Korbak, T. (2015) ‘Scaffolded Minds And The Evolution Of Content In Signaling Pathways’, *Studies in Logic, Grammar and Rhetoric* 41(1): 89–103.
- Kim, J. (1990) “Supervenience as a Philosophical Concept”, *Metaphilosophy*, 21: 1–27
- Kim, J. (1992) “Multiple Realization and the Metaphysics of Reduction”, *Philosophy and Phenomenological Research*, 52: 1–26.
- Langland-Hassen (2015) ‘Imaginative Attitudes’, *Philosophy and Phenomenological Research* 90(3): 664-686.
- Lettvin, J., Maturana, H., McCulloch, W. & Pitts, W. (1959) ‘What the Frog’s Eye Tells the Frog’s Brain’, *Proceedings of the IRE* 47(11): 1940-1951.
- Marr, D. (1982) *Vision*, San Francisco: Freeman.
- McGinn, C. (2004) *Mindsight: Image, Dream, Meaning*, Harvard University Press.
- Merleau-Ponty, M. (1962) *The Phenomenology of Perception*, trans. C. Smith, Routledge.
- Millikan, R., (1984) *Language, Thought and Other Biological Categories*, Cambridge, MA: MIT Press.
- Millikan, R. (1991) “Speaking Up for Darwin” in Loewer, B. & Rey, G. (eds.) *Meaning in Mind: Fodor and his critics*, Cambridge, MA: Blackwell,
- Millikan, R. (1993) *White Queen Psychology and Other Essays for Alice*, Cambridge, MA: MIT Press.
- Nagel, E. (1935) “The Logic of Reduction in the Sciences”, *Erkenntnis*, 5: 46–52

- Nagel, E. (1951) *The Structure of Science*, New York: Harcourt, Brace and World.
- Neander, K. (1995) “Malfunctioning and Misrepresenting”, *Philosophical Studies* 79: 109–141.
- Newell, A. & Simon, H. (1961) ‘Computer Simulation of Human Thinking’, *Science* 134: 2011-2017.
- O’Regan, K. & Noe, A. (2001) ‘A Sensorimotor Account of Vision and Visual Consciousness’, *Behavioral and Brain Sciences* 24: 939-1031
- Papineau, D. (1984) “Representation and Explanation”, *Philosophy of Science* 51: 550–72.
- Papineau, D. (1997) “Teleosemantics and Indeterminacy”, *Australasian Journal of Philosophy* 76: 1–14.
- Papineau, D. (2001) “The Status of Teleosemantics, or How to Stop Worrying about Swampman”, *Australasian Journal of Philosophy* 79: 279–89.
- Papineau, D. (2016) ‘Teleosemantics’, in D. L. Smith (ed.) *How Biology Shapes Philosophy: New Foundations for Naturalism*, Cambridge University Press.
- Pietroski, P. (1992) “Intentionality and Teleological Error”, *Pacific Philosophical Quarterly* 73 (3): 267-82.
- Price, C. (1998) “Determinate Functions”, *Noûs* 32: 54–75.
- Putnam, H. (1975) “The nature of mental states”, in *Philosophical Papers* (Volume 2), 429–440, Cambridge: Cambridge University Press.
- Ramsey, W. (2007) *Representation Reconsidered*, Cambridge University Press.
- Reboul, A. (2011) ‘Radical embodied cognition vs. "Classical" embodied neuroscience’, *The Journal of East China Normal University* 43 (6), 1-9.
- Rolla, G. (2017) “Contentless basic minds and perceptual knowledge”, *Unisinos Journal of Philosophy* 18(1): 47-56.
- Rosenberg, A. (2006) *Darwinian reductionism: or, how to stop worrying and love molecular biology*, Chicago: University of Chicago Press.
- Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*, Cambridge University Press.
- Turvey, M., Shaw, R., Reed, E. & Mace, W. (1981) ‘Ecological laws of perceiving and acting: in reply to Fodor and Pylyshyn’, *Cognition* 9: 237-304.
- Van Gelder, T. (1995) ‘What might cognition be if not computation?’ *Journal of Philosophy* 91: 345-381.
- Von Uexküll, J. (1934 [2011]) *A Foray into the Worlds of Animals and Humans, with A Theory of Meaning*, trans. Joseph D. O’Neil, University of Minnesota Press.
- Walter, W. G. (1950) ‘An Electromechanical Animal’, *Dialectica* 4(3): 206-213.
- Wiener, N. (1948) *Cybernetics, or Control and Communication in the Animal and in the Machine*, New York: Wiley.
- Wiese, W. (2017) ‘Action Is Enabled by Systematic Misrepresentations’, *Erkenntnis* 82 (6): 1233-1252.