<https://helda.helsinki.fi>

Making too many enemies : Hutto and Myin's attack on computationalism

Kuokkanen, Jesse

2018

Kuokkanen , J & Rusanen , A-M 2018 , ' Making too many enemies : Hutto and Myin's attack on computationalism ' , Philosophical explorations : an international journal for the philosophy of mind and action , vol. 21 , no. 2: Enactivism, Representationalism, and Predictive Processing , pp. 282-294 . <https://doi.org/10.1080/13869795.2018.1477980>

<http://hdl.handle.net/10138/248083>

<https://doi.org/10.1080/13869795.2018.1477980>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/325228919>

Making Too Many Enemies: Hutto and Myin's attack on computationalism

Article in *Philosophical Explorations* · May 2018

DOI: 10.1080/13869795.2018.1477980

CITATIONS

0

READS

200

2 authors:

Jesse Kuokkanen

University of Helsinki

1 PUBLICATION 0 CITATIONS

[SEE PROFILE](#)

Anna-Mari Rusanen

University of Helsinki

19 PUBLICATIONS 36 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:

The Many Faces of Relevance (<http://www.many-faces-of-relevance.org/>) [View project](#)

Making too many enemies: Hutto and Myin's attack on computationalism

Jesse Kuokkanen^{a*} and Anna-Mari Rusanen^b

^a*Social and Moral Philosophy, University of Helsinki, Helsinki, Finland;* ^b*Cognitive Science, Department of Digital Humanities, University of Helsinki, Helsinki, Finland*

We analyse Hutto and Myin's three arguments against computationalism (Hutto and Myin 2012, 2017; Hutto et al. forthcoming). *The Hard Problem of Content* targets computationalism that relies on semantic notion of computation, claiming that it cannot account for the natural origins of content. *The Intentionality Problem* is targeted against computationalism using non-semantic accounts of computation, arguing that it fails in explaining intentionality. *The Abstraction Problem* claims that causal interaction between concrete physical processes and abstract computational properties is problematic. We argue that these arguments are flawed and are not enough to rule out computationalism.

Keywords: radical enactive cognition; radical enactivism; computationalism; computational theory of mind; semantic computation; non-semantic computation

ISSN 0950-3110 print/ISSN 1473-348X online

© 2018 Taylor & Francis

DOI: 10.1080/XXXXXXXXXXXXXXXXXX

<http://www.informaworld.com> (live link)

1. Introduction

In recent years, many have argued that there can be cognitive science without the computational hypothesis of mind (see Chemero 2009; Chemero and Silberstein 2008; Stepp et al. 2011; Van Gelder 1995; Van Gelder and Port 1995; Varela, Thompson, and Rosch 1992). One such a proposal can be found in Hutto & Myin's Radical Enactive Cognition (REC) (2012, 2017; Hutto et al. forthcoming)¹.

According to REC, cognition is, using Hutto and Myin's characterization, a multi-storey story: there is both non-representational basic cognition and socio-culturally scaffolded, potentially representational cognition. With this distinction Hutto and Myin want to contrast their account with a view they dub as contemporary cognitivism, which 'takes it to be axiomatic that "the mind represents and computes" (Branquinho 2001, xv)' (Hutto and Myin 2017, 3). In this picture, cognitivism consists of two 'twin pillars', computationalism and representationalism, and is 'methodologically committed to providing explanations of a mechanistic variety'. According to Hutto and Myin, this approach is a 'single-storey' account of cognition because it, according to them, suggests that *all* cognition is representational. Hutto and Myin aim to reject this kind of cognitivism which takes cognition as computation of inner representations.

However, their earlier arguments are targeted mainly against representationalism and not against computationalism per se. Instead of proper and careful arguments, they have made some short but explicit remarks, such as 'representation and computation ... are not definitive of, and do not form the basis of, all mentality' (Hutto and Myin 2012, 3), and that their '*rejection of ... the twin representational and computational pillars of cognitivism, is motivated by the avoidance of deep theoretical mysteries*' (Hutto and Myin 2017, 51, italics added). Only in their most recent paper (Hutto et al. forthcoming) they address the issue of computationalism in a more detail.

¹ There has already been discussion on the arguments against computationalism in classical enactivism (Dewhurst and Villalobos 2017; Villalobos and Dewhurst 2017, 2018).

In this article, we argue that all Hutto and Myin's arguments against computationalism fail. Hence, they have not provided sufficient reasons for rejecting computationalism. However, we emphasize that (i) we do not intend to argue for or against computationalism, (ii) we do not intend to argue for or against enactivism, and (iii) we do not intend to argue that computationalism is compatible with REC.

This paper is organised as follows: In section 2 we briefly describe what computationalism entails. Section 3 presents Hutto and Myin's first argument (2012) against computationalism, the Hard Problem of Content (HPC). We argue that it leaves non-semantic views of computation untouched. Because of this, even if one agrees with the HPC, Hutto and Myin's early rejection of computationalism *per se* fails. In section 4 we look at an argument we call the Intentionality Problem (Hutto and Myin 2017), which comes in the form of a dilemma. Unlike their earlier objection, this argument is aimed also against non-semantic accounts of computation. We argue that the Intentionality Problem is flawed since it rests on a confusion concerning what computationalism entails. In section 5 we present their most recent objection which we call the Abstraction Problem (Hutto et al. forthcoming), and argue that it is based on a confusion concerning the notion of abstraction. Our conclusion (section 6) is that Hutto and Myin do not provide sufficient reasons for rejecting computationalism, and that further arguments are needed if they wish to make such a move.

2. What is computationalism?

To cut it short, computationalism is a view according to which cognition is computational. However, it comes many forms, and certain distinctions should be made to respect the complexity at hand.

First, computation may be studied mathematically by formally defining computational objects such as algorithms and proving theorems about their properties. While this kind of study of computational objects and their properties belongs to mathematics, physical computation tries to explain what it means for a physical system to perform a computation. From this perspective, computationalism is about physical

computation: assuming that cognition is a physical phenomenon, does it compute-- and if it does, why? (Piccinini 2007b, 2015; Sprevak 2010) Second, computationalism can be defended in stronger and weaker versions. The weaker versions take only some cognitive processes to be computational, while the strongest formulations take all cognitive processes to be computational.

The third distinction concerns the nature of computation itself: there are semantic and non-semantic views of computation. For example, one may support a semantic view according to which it is necessary for computation to be manipulation of contentful representations (e.g., Fodor 1975), or a non-semantic view which denies this (e.g., Piccinini 2015; Stich 1983). There is also another sense in which an account can be semantic. This concerns the question of how computations are to be individuated. According to the semantic accounts, we need representations or semantic content to individuate computations (e.g., Fodor 1975, 1981). Non-semantic accounts of computational individuation deny this, claiming that semantics is not needed to individuate computations. For example, according to Piccinini (2007a, 2015), computational states are individuated by their functional properties, and their functional properties are specified through an appeal to a mechanism in a way that need not refer to any semantic properties.² In this paper, by a semantic account of computation we refer to an account that is semantic in either of the ways described above.

Fourth, there is a distinction between ontological and explanatory computationalism. The ontological computationalists think that cognitive and neurocognitive systems literally perform computations, while the explanatory computationalists see it as useful to ascribe computations to cognition without making any explicit claims about whether those processes really are computational. In contrast to some famous intuitions (Fodor 1975), computationalism is not tied to the view that cognition should - or could - be explained *only* computationally. Rather, computationalism entails merely that cognitive processes allow for a computational explanation.³ This does

² In reality, the question of semantics of computation is even more messy than described here. One can talk about semantics at least at the level of explanation (algorithms versus computations), characteristics of representations, individuation of computations, and implementation.

³ It is also worth noting that there is not consensus on the meaning of computational explanation itself (Rusanen and Lappi 2007).

not imply that cognition or cognitive phenomena *must* be explained computationally, or that computational explanations would be the *only* acceptable type of explanations for cognitive phenomena. Instead, many defend the view that full explanations of implemented cognitive phenomena may require also neural and other type of explanations⁴.

From now on, by computationalism we mean a mainstream view of computational theory of cognition, according to which cognitive states and cognitive processes are computational. The take-home message of this section is that computationalism does not entail any specific sort of explanation nor a way of individuating computations. It is a claim concerning the underlying nature of cognitive processes.

3. Computationalism and the Hard Problem of Content

In their first book, Hutto and Myin (2012) do not explicate what they mean by computation or computationalism, but they explicitly claim that ‘representation and computation ... are not definitive of, and do not form the basis of, all mentality’ (Hutto and Myin 2012, 3). A major part of the argumentation in the book is devoted to the Hard Problem of Content (HPC)⁵.

The HPC is targeted primarily against traditional representationalist accounts of cognition, which take representation and content to be defining features of cognition. By content, Hutto and Myin mean the semantic features of representations. My utterance ‘there is a mug on the table’ has content, since it has conditions of satisfaction: it is true in the case that there is a mug on the table, and it is false in the case there is not. Conditions of satisfaction form the minimum requirement for something counting as a representation.⁶ Thus, if something is a representation, it

⁴ For discussion, see (Egan 2017; Kaplan 2011; Miłkowski 2013; Piccinini and Craver 2011; Rusanen and Lappi 2016; Shagrir 2010; Weiskopf 2011).

⁵ Whether the HPC should be seen as a genuine worry for representationalism has already received a lot of discussion. See (Colombo 2014; Miłkowski 2015; Shapiro 2014) for critical remarks and (Hutto and Myin 2014, 2017) for replies.

⁶ There are also proposals that representations need not involve such content. According to Hutto and Myin, however, this kind of talk muddles the conversation. They suggest that if one is not willing to talk about content, it is better to abandon the notion of representation altogether.

has content. According to the HPC, traditional semantic theories of cognition cannot give a scientifically respectable story of content and hence, we should abandon the idea that cognition always involves contentful representations.

According to Hutto and Myin, all naturalists⁷ who posit content to cognition must somehow answer the HPC. Their own answer to the HPC is a neo-pragmatist, socio-cultural story of the origins of content. According to REC, “the primary bearers of content are semantically articulated symbols, occurring in appropriate dynamic patterns” (Haugeland 1990, 412)’ (Hutto and Myin 2017, 124). This story claims that content appears on the scene only at the ‘socio-culturally scaffolded level’ of cognition and hence, there is no content at the level of basic cognition. This is bad news for representationalists who claim, according to Hutto and Myin, that a *defining feature* of cognition is the manipulation of representations. If representationalists accept that there is cognitive behaviour also below the so-called socio-cultural level of cognition, this entails that there are representations also at the basic level. However, representations at the level of basic cognition are problematic, since representations are entities that have semantic features such as sense, content and reference, and those are features that can be found only at the socio-cultural level of cognition. If one wants to remain naturalistic, such content needs a scientifically respectable explanation which, according to Hutto and Myin, is not available at the level of basic cognition. They analyse a teleosemantic answer to the HPC, which they consider as the most promising and popular one, but argue that it cannot solve the problem (Hutto & Myin 2012, ch. 4, 2017, 41-45). Their conclusion is that we should not posit representations to basic cognition.

The HPC can be seen as a problem for certain kind of computationalism, since some (e.g., most of the traditional) semantic accounts of computation either take computation to be manipulation of representations or claim that computations must be individuated semantically. In either case, representations are involved in both basic and socio-cultural cognition which, according to Hutto and Myin, makes

⁷ Hutto and Myin do not give a precise definition of what they mean by naturalism. Roughly, for them naturalism amounts to giving a scientifically respectable story of a certain phenomenon. They distinguish strict naturalism from a more relaxed naturalism. When the strict naturalism allows only the so-called hard sciences to appear in a naturalist story, the more relaxed naturalism takes on board more explanatory power by recruiting ‘the full range of scientifically respectable resources, drawing on the findings of a wide variety of sciences that include not just the hard sciences’ (2017, 124)

computationalism unable to answer the HPC. However, as argued in the section 2, not all accounts of computation take computation to entail representations. For this reason, non-semantic and semantically neutral accounts of computation are not affected by the HPC⁸. This leads to the conclusion that Hutto and Myin's early rejection of computationalism per se is unjustified.

4. The Intentionality Problem

The objection we call Intentionality Problem (Hutto and Myin 2017) is a more explicit argument against computationalism. It is a part of a dilemma argument, which is supposed to affect also non-semantic and semantically neutral accounts of computation. The first horn of the argument is the Hard Problem of Content (HPC): the traditional account of cognitivism, which takes computationalism and representationalism as the twin pillars of cognition, is unable to answer the HPC. This can be read also in a way that any form of computationalism that both relies on a semantic notion of computation and is strong enough to claim that cognitive processes at the level of basic cognition are computational is unable to answer the HPC. As we have already argued above, the HPC does not concern computationalism per se but only certain kinds of semantic computationalism.⁹

If one wants to make such a move to avoid the HPC, that is, drop out representations out of the picture and adopt a non-semantic account of computation, she is faced with the second horn of the dilemma, which is the Intentionality Problem. The motivation behind the Intentionality Problem is that the standard approach in explaining intentionality is, according to Hutto and Myin, to use representations: representations are entities that have semantic features such as reference, which are handy for the task. Views which hold that cognition is entirely a matter of non-representational computations and offer no successor notion for

⁸ To be more specific, only sufficiently strong variants of semantic computationalism are affected by this argument: a weak variant semantic computationalism might be seen unaffected by the HPC if it does not claim that cognitive processes at the level of basic cognition are computational.

⁹ Whether the HPC should be seen as a genuine worry for representationalism has already received a lot of discussion. See (Colombo 2014; Miłkowski 2015; Shapiro 2014), and (Hutto and Myin 2014, 2017) for replies.

representation 'will be unable to explain how organisms relate to and connect with targeted aspects of their worldly environments' (Hutto and Myin 2017, 50). In a nutshell, if one accepts the conclusion of HPC and makes the basic level of cognition representation-free, computation alone cannot explain intentionality, according to Hutto and Myin. In other words, to solve the dilemma one must either provide an answer to the HPC or introduce a successor notion for semantic representations that can explain intentionality.

Hutto and Myin consider a possible reply to the Intentionality Problem, which is to come up with a successor notion for representation that can account for intentionality. The candidate notion they consider is semantic sensitivity, by which they mean a property of computations that Piccinini¹⁰ (2004, 2006, 2007a, 2015, 32) and Rescorla (2012, 2014) consider: even if computations are not representational in nature and need not be identified by representations, computations can still be sensitive to semantic properties. This, however, makes the view again vulnerable to the HPC: 'We have been at pains to show that paying for that assumption requires facing up to the HPC' (Hutto and Myin 2017, 51). This all leads to the conclusion that computationalism is either tied to the HPC or cannot explain intentionality.

Here is our reconstruction of the argument. First, Hutto and Myin asks us to agree on two things:

- (i) All cognition is not alike. There is basic cognition which is not representational, and then there is socio-culturally scaffolded cognition which can be representational.
- (ii) Even though basic cognition is not representational, it can be intentional.

Then, the argument continues:

- (iii) The traditional way of dealing with intentionality is positing contentful representations to cognition.
- (iv) Using contentful representations in explaining intentionality makes the HPC unsolvable, since it posits representations to all levels of cognition.

¹⁰ For the traditional discussion, see (Fodor 1981).

- (v) If one tries to solve the HPC by letting go of representations and adopt a non-semantic or semantically neutral account of computation, she is unable to explain intentionality without a successor notion. Thus,
- (vi) if representations are left out, computationalism needs the notion of semantic sensitivity to account for intentionality. However,
- (vii) semantic sensitivity leads back to the HPC, which computationalism cannot solve.
- (viii) Since computationalism cannot solve the HPC, computationalism is rejected.

We argue that even if we accepted, for the sake of the argument, premises (i) -- (v), the crucial premise is (vi) false¹¹. Computationalism does not entail a specific story of intentionality and hence, it does not need to rely on semantic sensitivity. For this reason, the argument is unsound.

Apparently, the reason why Hutto and Myin assume that computationalism, which relies on non-semantic or content-neutral computation, needs the notion of semantic sensitivity is the fact that Piccinini and Rescorla talk about semantic sensitivity when formulating their stances on computation. Hutto and Myin (2017) quote Piccinini (2015), drawing a conclusion that Piccinini's theory of computation is tied to semantic sensitivity:

They allow that even if computations are not essentially individuated by semantic properties—even if computations have a wholly nonsemantic and mechanistic nature—they can still be sensitive to semantic properties (Rescorla 2012a, 2014; Piccinini 2015). Why so? The reason this is the preferred view is clear enough. Such theorists feel compelled to assume that “there are ... semantic properties that relate many computing systems to their environment [...]” (Piccinini 2015, 32). We have been at pains to show that paying for that assumption requires facing up to the HPC [Hard Problem of Content] in one of the ways described above. (Hutto and Myin, 2017)

¹¹ As mentioned earlier in the fn. 6, one can also try to argue that the HPC itself is not a proper problem. In this paper, we try another strategy.

However, it is problematic to infer on that computationalism relying on non-semantic or semantically neutral account of computation *needs* the notion of semantic sensitivity. This is not the case. First, even though Piccinini is vague and sketchy when it comes to semantic sensitivity¹², he seems to be talking about internal semantics. Internal semantics is ‘fully determined by the functional and structural properties ... *independently of any external semantics*’ (Piccinini 2015, 135, italics added). In contrast, philosophers are traditionally talking about external semantics, that ‘relates a state to things other than its computational effects ... including objects and properties in the external world’ (Piccinini 2015, 135). Piccinini makes it clear that in order for computations to be sensitive to semantic properties, there has to be a causal mechanism between them, meaning that computations cannot be sensitive to “wide” meanings.¹³ Putting this another way, the way Piccinini defines computation does not force him, or any other computationalist, to accept the idea of semantic sensitivity in a way that Hutto and Myin envisions.

Second, Piccinini does not even try to explain intentionality by appealing to computations alone. As Piccinini explicitly writes, computation may or may not contribute in explaining intentionality: ‘whether computation has a semantic nature should not be confused with whether computation ... explains ... intentionality ... I will remain neutral on whether being computational contributes to explain original intentionality’ (Piccinini 2015, 32, italics added). The nature of intentionality is a separate--and perhaps partially empirical--question, and so is the question whether intentionality turns out to be computational or not. In sum, computationalism does not entail anything about the appropriate strategy for explaining intentionality. However, depending on how strong a thesis one supports, naturalist computationalism must be

¹² In the very same passage that Hutto and Myin use, Piccinini writes that ‘[t]he issue of semantic properties and their role in computational causation is too complex a topic to address it here’ (Piccinini 2015, 32). Moreover, the meaning of semantics in the passage is not clear, and Piccinini writes that ‘whether computation is sensitive to semantic properties depends on what “sensitive” and “semantic property” means’ (Piccinini 2015, 32).

¹³ From this perspective, when Piccinini is talking about semantic sensitivity, the meaning of semantics is such that it seems compatible even with Hutto and Myin’s ideas of informational sensitivity, contentless signs and biological functions.

compatible with the best scientific theory of intentionality. And still, importantly, it is not up to computationalism to come up with an explanation of intentionality.¹⁴

Because of this, there seems to be no a priori reason why computationalism could not be seen compatible with Hutto and Myin's own account of intentionality, according to which one should 'seek to explain how full-blown cognitive ... capacities might have been built-up from simpler, less abstract aspects' (Muller 2014, 169) and 'surrender the idea that basic forms of intentionality need involve correctness or satisfactory conditions of any kind' (Hutto and Myin 2017, 101). They aim to give a naturalistic account of "basic intentionality" with scientifically respectable notions, such of biological functions, dispositions, contentless signs, and informational sensitivity, to mention a few. They continue by explicating that 'This is more or less the REC view: in basic kinds of cognition an organism's skillful engagements with the world are best understood in embodied, enactive, and nonrepresentational ways' (101-102). However, Hutto and Myin do not provide any reasons for why these should not be seen or described as computational. In fact, some of the most elegant neurocognitive theories of "skillful engagements" can be, and are, given in purely computational terms (e.g., Franklin and Wolpert 2011; Wolpert, Doya, and Kawato 2003; Wolpert and Ghahramani 2000).

To sum up, in their 2017 book, Hutto and Myin do not consider the possibility that computationalism could adopt their own story of intentionality. Instead, they assume that computationalism is forced to come up with a "single-storey story" of cognition, which entails that computationalism is bound to be unable to answer the HPC. According to their treatment, the options open for computationalism are either to posit semantics to every level of cognition, or refrain from postulating semantics to cognition altogether. However, as we have argued, if a computationalist adopts a non-semantic account of computation it does not imply that she is forced to lean on a semantic notion to explain intentionality, even if some computationalists do so. Hence, we conclude, the Intentionality Problem is unsound.

¹⁴ Elsewhere, Morgan and Piccinini (2018) defend a multi-level mechanistic account of intentionality that is committed to the Representational Hypothesis of Cognition. Despite this fact, the explanation of intentionality and the role of semantics is, in principle, independent of the computational framework.

Apparently, Hutto and Myin have noticed the shortcomings of their arguments, since in their forthcoming paper they do not outright reject computationalism. Instead, they present a new objection that is targeted especially against non-semantic account of computation which, according to them, makes computationalism relying on such an account “problematic”.

5. The Abstraction Problem

In their forthcoming paper, Hutto et al. they do not reject computationalism altogether, but ‘question the ability of ... [non-representational] theories [of computation] to deliver a metaphysically robust, naturalistic account of computation of the sort needed to support [computationalism]’. We call this new objection the Abstraction Problem^{15,16}.

According to the Abstraction Problem, the main challenge for the advocates of neomechanistic, non-semantic account of computation one like Piccinini’s is to explain ‘how concrete neural processes could causally manipulate abstract, medium independent vehicles’ (Hutto et al., forthcoming). Since, according to Hutto et al., defenders of the mechanistic theory of computation offer ‘no account of how such manipulations might be achieved’, they, Hutto et al. claim, cannot explain the computational basis of cognition. However, as we see it, their argument rests on a misunderstanding concerning the notion of abstraction and hence, it fails.

Let’s take a look at the Abstraction Problem in a more detail. According to Hutto et al., the ‘apparent obstacle’ in defending computationalism through an appeal to a

¹⁵ Instead of using the term computationalism, Hutto et al. talk about “Computational Basis of Cognition” (CBC). However, it seems to us that CBC is just another name for computationalism.

¹⁶ Three remarks are relevant concerning the notion of computationalism in their forthcoming paper. First, Hutto et al. are talking about the strongest possible variant of computationalism: all cognitive processes must be computational. Hence, they fail to do justice to all the different variants of computational thinking. Also, according to them, computationalism maintains that ‘computation is a, if not the, explanatory basis for cognition’. The notion of “explanatory basis” is not clarified in the paper, which makes it ambiguous. Third, they make an implicit claim that computationalism is committed to internalism, that is, to the thesis that we should explain cognitive processes by looking at the brain-based neural processes. This is false: computationalism can be seen compatible with wide computationalism, which denies internalism (Kersten 2016; Piccinini 1994, 2004, 2015; Wilson 1994, 2004).

mechanistic theory of computation, is the nature of neural processes. As Hutto et al. put it, there are ‘clear dissimilarities between what happens in brains and what happens in artefactual computers’ (Hutto et al., forthcoming). Hutto et al. refer to Piccinini and Bahar’s article (2013)¹⁷, in which they analyse neural phenomena, such as the so-called spike trains, as the primary candidates ‘for interneural long-distance signaling’ (Piccinini and Bahar 2013, 462). However, as Piccinini and Bahar remark, ‘typical neural signals, such as spike trains ... are neither continuous signals nor strings of digits’ (453). Hence, spike trains cannot be analyzed through classical computability theory (Piccinini 2008; Piccinini and Bahar 2013).

Instead of drawing the conclusion that these type of neural signals are not computational, Piccinini and Bahar suggest that they are computational *sui generis*.¹⁸ Hutto et al. claim that one would be ‘equally justified in concluding that brains do not compute’. However, this is not the case. Piccinini and Bahar do not make merely an ad hoc proposal that since spike trains--or other similar neural processes--are not computational in the standard--that is, digital--sense, they are computational in some other sense.

Instead, that conclusion follows from an interpretation of empirical evidence through the lense of Piccinini’s theory of physical computation. In Piccinini’s account, digital and analog computation are subclasses of generic computation. Even if spike trains are neither digital nor analog computation, they do fall within the scope and criteria of generic computation mainly because they seem to be, according to Piccinini and Bahar, medium independent. Hence, while Piccinini and Bahar provide an argument for the claim that spike trains should be seen as a form of computation, Hutto et al. provide none for their “equally justified”, alternative and contrary conclusion.

Piccinini and Bahar note that spike trains seem to be medium independent since, according to the current evidence, ‘functionally relevant aspects of neural processes

¹⁷ For the original argument, see (Piccinini 2008).

¹⁸ Piccinini and Bahar give two arguments for generic computationalism. The first is ‘*the argument from the functional organization of the nervous system*’, which is the one we summarized. The second one is ‘*the argument from semantic information processing*’. In that argument, Piccinini and Bahar claim that at least some cognitive processes, like language processing, seem to require representations in Hutto & Myin’s sense. However, Piccinini and Bahar’s argument for generic computationalism does not hang on this claim.

depend on dynamical aspects of the vehicles--most relevantly, spike rates and spike timing ... and may be implemented either by neural tissue or by some other physical medium, such as a silicon-based circuit ... thereby qualifying as proper vehicles for generic computation' (Piccinini and Bahar 2013, 462). In other words, these 'vehicles' can be seen as medium independent because they may be implemented by other appropriate physical medium. This, Piccinini and Bahar claim, is the main reason why spike trains qualify as vehicles for computation.

However, according to Hutto et al., '[i]t is questionable ... that the neural ... processes ... have the feature of being medium independent' (Hutto et al. forthcoming). It is questionable, because '[t]here is reason to doubt that neural events could contribute to cognitive work if that work really requires the concrete manipulation of medium-independent vehicles.' Furthermore, '[t]he trouble is that if medium independent vehicles are defined by their abstract properties then it is unclear how such vehicles could be concretely manipulated' and hence, computationalists have 'no conception of how' these vehicles could be processed neurally. Thus, according to the Abstraction Problem, the real difficulty is to explain how concrete neural processes 'could causally manipulate abstract, medium independent vehicles'.

As we see it, the Abstraction Problem is based on a misunderstanding concerning what the abstract nature of medium independent vehicles means. In one sense, Hutto et al. confuse the metaphysical and explanatory claims (see section 2 for the distinction). Namely, for Piccinini, concrete computations and their vehicles can be *defined* independently of the physical media that implement them (Piccinini 2015 120-122). Putting this another way, for Piccinini the abstraction is a matter of omitting irrelevant features, and not a matter of assuming that these features exist in a metaphysically abstract way. For example, if one makes a model of the famous aeroplane, "the Spirit of St.Louis", one may omit some of the target system's details and that way construct an abstracted model of the original plane. This does not affect the properties of the original plane. Or, if one wants to explain the aerodynamics of the Spirit of St.Louis, one may ignore many of the plane's properties, such as its color, as irrelevant for the explanation, and thus have an abstracted explanation of a target system (Craver 2006; Craver and Kaplan 2018). In

the same sense, for Piccinini, the medium independent vehicles are abstract in a sense that they can be described at different levels of abstractions.

In other words, to describe concrete computations and the vehicles they manipulate one need not consider all of their specific physical properties (Piccinini 2015, 120-122). Instead, one may consider only the properties that are relevant to the computation according to the rules that define the computation. Importantly, these higher-level properties, such as the medium independent properties, can be realized in different lower-level properties that constitute different mechanisms at the immediate lower mechanistic levels. This does not make the higher-level properties any less metaphysically concrete.

However, Hutto et al. seem to think that when talking about the description of medium independent vehicles, Piccinini uses the notion of abstract in a way that makes the vehicles, not the descriptions of them, metaphysically abstract. Namely, Hutto et al.'s claim about the neo-mechanists' difficulties with 'explaining how... abstract entities can be causally manipulated' by 'concrete...processes' would not make any sense if it was not committed to the metaphysical interpretation of abstract vehicles.

But, the claim that computational descriptions are committed to the existence of abstract medium independent vehicles is a result of a confusion between the abstractness as a feature of models or descriptions and of 'the processes being modelled' (Polger and Shapiro 2016, 166). Instead, as Polger and Shapiro emphasize, the medium independence of vehicles should be interpreted as a claim that the computational models describe the target phenomena in an abstract way rather than as a claim that the properties or processes of target systems can exist in an abstract way (Polger and Shapiro 2016, 166)¹⁹. Namely, one can have descriptions of concrete vehicles as medium independent and abstract, but one

¹⁹ According to Polger & Shapiro (2016), also Piccinini and Bahar (2013) and Piccinini (2015) make this mistake. However, Polger & Shapiro (2016) focuses mostly on the multiple realizability of medium independent vehicles, and it is not obviously clear to what extent that argument can be utilized in a way that Hutto et al. (forthcoming) do. Moreover, even if there were certain problems in Piccinini's account, that issue is beyond the scope of this paper.

cannot have (metaphysically) concrete vehicles that are (metaphysically) abstract at the same time.

6. Conclusions

In this article, we have analysed the objections against computationalism presented by radical enactivism (Hutto and Myin 2012, 2017; Hutto et al. forthcoming). Hutto and Myin's early rejection of computationalism (Hutto and Myin 2012) is based on the Hard Problem of Content (HPC), according to which traditional views, such as teleosemantic accounts that posit representations to all levels of cognition, are unable to give a naturalist explanation of representations at the basic level. However, the HPC considers only certain semantic accounts of computation, making Hutto and Myin's rejection of computationalism *per se* unjustified.

According to their more interesting dilemma argument (Hutto and Myin 2017), computationalism is unable either to answer the HPC or to explain intentionality. The first horn of the argument is the HPC. If one tries to avoid the HPC by leaving representations out of the picture, she is faced with the second horn of the argument, which is the Intentionality Problem. According to the Intentionality Problem, positing representations is the standard strategy for explaining intentionality, and if one abandons representations, she cannot explain intentionality without a successor notion. Hutto and Myin consider the notion of semantic sensitivity as a possible reply and successor notion for the task, and argue that it does not resolve the issue, since it makes computationalism again semantic and thus, unable to answer the HPC. We have argued that the second horn of the dilemma is unsound. Hutto and Myin's suggestion that computationalism needs the notion of semantic sensitivity for explaining intentionality is false. The main problem of their argument is that the way Hutto and Myin treat computationalism does not give computationalism any other option but to face the HPC: the only successor notion they consider for explaining intentionality is semantic in nature which, by definition, makes computationalism unable to answer the HPC. However, if a computationalist adopts a non-semantic account of computation, it does not follow that she is forced to lean on a semantic account to explain intentionality, even if some computationalists do so. Importantly,

Hutto and Myin do not consider the possibility that computation could be seen compatible with their own story of intentionality.

Their recent objection to computationalism is a problem which we call the Abstraction Problem (Hutto et al. forthcoming). According to the Abstraction Problem, the main challenge for the advocates of neomechanistic, non-semantic account of computation is to explain 'how concrete neural processes could causally manipulate abstract, medium independent vehicles' (Hutto et al. forthcoming). We have argued that the Abstraction Problem rests on a misunderstanding concerning the notion of abstraction and hence, it fails.

Our diagnosis is that the main problem behind all three arguments is a confusion concerning what computationalism entails: Computationalism is a thesis about the underlying nature of cognitive processes, and it comes in many forms. Moreover, computationalism by itself does not entail commitments to representations, any specific story of intentionality, or metaphysically abstract entities.

Acknowledgements

- (1) This article is based on a presentation given at the Evolving Enactivism workshop in Bochum, 5 December 2017, and on Jesse Kuokkanen's master's thesis (2018).
- (2) We thank Gualtiero Piccinini and two autonomous referees for their valuable comments on this article.

Notes on contributors

Mr. Jesse Kuokkanen is working as a research assistant in TINT – Centre for Philosophy of Social sciences, and finishing his master's degree in Theoretical Philosophy at University of Helsinki.

Dr. Anna-Mari Rusanen is a university lecturer in Cognitive Science, Department of Digital Humanities, at University of Helsinki.

References

- Branquinho, J. 2001. *The Foundations of Cognitive Science*. Oxford University Press.
- Chemero, A. 2009. *Radical Embodied Cognitive Science*. MIT Press.
- Chemero, A., and M. Silberstein. 2008. "After the Philosophy of Mind: Replacing Scholasticism with Science." *Philosophy of Science*, 75 (1): 1–27.
- Colombo, M. 2014. "Neural Representationalism, the Hard Problem of Content and Vitiating Verdicts: A Reply to Hutto & Myin." *Phenomenology and the Cognitive Sciences* 13 (2): 257–274.
- Craver, C. F. 2006. "When Mechanistic Models Explain." *Synthese* 153 (3): 355–76.
- Craver, C. F. 2007. *Explaining the Brain: What a Science of the Mind-Brain could be*. Oxford University Press.
- Craver, C.F., and D. M. Kaplan. 2018. "Are More Details Better? On the Norms of Completeness for Mechanistic Explanations." *The British Journal for the Philosophy of Science*. <https://doi.org/10.1093/bjps/axy015>
- Dewhurst, J., and M. Villalobos. 2017. "The Enactive Automaton as a Computing Mechanism." *Thought: A Journal of Philosophy* 6: 185–192.
- Dreyfus, H. L. 2002. "Intelligence without representation: Merleau-Ponty's critique of mental representation." *Phenomenology and the Cognitive Sciences* 1: 367–383.
- Egan, F. 2017. "Function-Theoretic Explanation and the Search for Neural Mechanisms." In *Explanation and Integration in Mind and Brain Science*, edited by David M. Kaplan, 145–163. Oxford University Press.
- Fodor, J. A. 1981. *Representations: Philosophical Essays on the Foundations of Cognitive Science*. MIT Press.
- Fodor, J. A. 1995. *The Language of Thought*. Harvard University Press.
- Franklin, D. W., and D. M. Wolpert. 2011. "Computational mechanisms of sensorimotor control." *Neuron* 72: 425–442.
- Haugeland, J. 1990. "The intentionality all-stars." *Philosophical Perspectives* 4: 383–427.
- Hutto, D., and E. Myin. 2012. *Radicalizing Enactivism: Basic Minds Without Content*. MIT Press.
- Hutto, D., and E. Myin. 2014. "Neural Representations Not Needed - No More Pleas, Please." *Phenomenology and the Cognitive Sciences* 13 (2): 241–256.
- Hutto, D., and E. Myin. 2017. *Evolving Enactivism: Basic Minds Meet Content*. MIT Press.

Hutto, D., E. Myin, A. Peeters, F. Zahnoun. Forthcoming. "The Cognitive Basis of Computation: Putting Computation In Its Place." In *The Routledge Handbook of The Computational Mind*, edited by M. Colombo, and M. Sprevak. London: Routledge.

Kuokkanen, J. (2018). "Radikaalin enaktivismin kritiikki komputationalismia vastaan." MA diss., University of Helsinki.

Miłkowski, M. 2013. *Explaining the Computational Mind*. MIT Press.

Miłkowski, M. 2015. "The Hard Problem of Content: Solved (Long Ago)." *Studies in Logic, Grammar and Rhetoric* 41 (1): 73-88.

Muller, H. D. 2014. "Naturalism and intentionality." In *Contemporary Philosophical Naturalism and Its Implications*, edited by B. Bashour, and H. D. Muller, 155-181. London: Routledge.

Morgan, A., and G. Piccinini. 2018. "Towards a Cognitive Neuroscience of Intentionality." *Minds and Machines* 28 (1): 119-139.

Piccinini, G. 2004. "Functionalism, Computationalism and Mental Contents." *Canadian Journal of Philosophy* 34: 375-410.

Piccinini, G. 2006. "Computational Explanation in Neuroscience." *Synthese* 153: 343-353.

Piccinini, G. 2007a. "Computational Explanation and Mechanistic Explanation of Mind." In *Cartographies of the Mind: Philosophy and Psychology in Intersection*, edited by M. Marraffa, M. DeCaro, and F. Ferretti, 23-36. Dordrecht: Kluwer.

Piccinini, G. 2007b. "Computing Mechanisms." *Philosophy of Science* 74: 501-526.

Piccinini, G., and S. Bahar. 2013. "Neural computation and the computational theory of cognition." *Cognitive Science* 37 (3): 453-488.

Piccinini, G. 2015. *Physical Computation: A Mechanistic Account*. Oxford University Press.

Polger, T.W., and L. A. Shapiro. 2016. *The multiple realization book*. Oxford University Press.

Rescorla, M. 2012. "Are Computational Transitions Sensitive to Semantics?" *Australasian Journal of Philosophy* 90 (4): 703-21.

Rescorla, M. 2014. "The Causal Relevance of Content to Computation." *Philosophy and Phenomenological Research* 88 (1): 173-208.

Rusanen, A-M., and O. Lappi. 2007. "The Limits of Mechanistic Explanation in Cognitive Science." In *Proceedings of the European Cognitive Science Conference 2007*, edited by S. Vosniadou, D. Kayser, and A. Protopapas, 284-289. Howe: Lawrence Erlbaum Associates.

Rusanen, A-M., and O. Lappi. 2016. "On Computational Explanations." *Synthese* 193 (12): 3931-3949.

Shapiro, L. A. 2014. "Radicalizing Enactivism: Basic Minds without Content." Review of *Radicalizing Enactivism: Basic Minds without Content*, by D. Hutto, and E. Myin. *Mind* 123 (489): 213-220.

Sprevak, M. 2010. "Computation, Individuation, and the Received View of Implementation." *Studies in the History and Philosophy of Science* 13 (2): 107-43.

- Stepp, N., A. Chemero, and M. T. Turvey. 2011. "Philosophy for the Rest of Cognitive Science." *Topics in Cognitive Science* 3: 425–437.
- Van Gelder, T. 1995. "What might cognition be if not computation?" *Journal of Philosophy* 92 (7): 345-81.
- Van Gelder, T., and R. F. Port, eds. 1995. *Mind as Motion*. MIT Press.
- Villalobos, M., and J. Dewhurst. 2017. "Why post-cognitivism does not (necessarily) imply anti-computationalism." *Adaptive Behaviour* 25 (3): 117–128.
- Villalobos, M., and J. Dewhurst. 2018. "Enactive autonomy in computational systems." *Synthese* 195 (5): 1891-1908.
- Weiskopf, D. 2004. "The Place of Time in Cognition." *British Journal for the Philosophy of Science* 55: 87–105.
- Weiskopf, D. 2011. "Models and mechanisms in psychological explanation." *Synthese* 183: 313-38.
- Wolpert, D. M., and Z. Ghahramani. 2000. "Computational principles of movement neuroscience." *Nature Neuroscience* 3: 1212–1217.
- Wolpert, D. M., K. Doya, and M. Kawato. 2003. "A unifying computational framework for motor control and social interaction." *Philosophical Transactions of the Royal Society B: Biological Sciences* 358 (1431): 593–602.
<http://doi.org/10.1098/rstb.2002.1238>