Against Computational Perspectivalism

Dimitri Coelho Mollo*

Forthcoming in *The British Journal for the Philosophy of Science*
Please only quote and cite the version of reference at
https://doi.org/10.1093/bjps/axz036

Abstract

Computational perspectivalism has been recently proposed as an alternative to mainstream accounts of physical computation, and especially to the teleologically-based mechanistic view. It takes physical computation to be partly dependent on explanatory perspectives, and eschews appeal to teleology in helping individuate computational systems. I assess several varieties of computational perspectivalism, showing that they either collapse into existing non-perspectival views; or end up with unsatisfactory or implausible accounts of physical computation. Computational perspectivalism fails therefore to be a compelling alternative to perspective-independent theories of computation in physical systems. I conclude that a teleologically-based, non-perspectival mechanistic account of physical computation is to be preferred.

1 Introduction

The notion of computation plays a central role in several scientific endeavours, and especially in the computer and cognitive sciences. These two fields — understood broadly — not only use computational tools and models in their investigations, but have as their subject matter physical systems believed to be computational in nature. This claim is somewhat more controversial in the case of the cognitive sciences, as the computational theory

*Science of Intelligence Cluster (SCIoI) & Berlin School of Mind and Brain & Institut für Philosophie, Humboldt-Universität zu Berlin.

of cognition, despite its widespread adoption since the Cognitive Revolution in the '50s, has found many detractors through the years. However, occasional scepticism notwithstanding, a large consensus has formed around the idea that cognitive systems are computational — a consensus that survives and thrives to this day (Kriegeskorte & Douglas 2018). Their computational nature, moreover, is supposed to help explain their cognitive nature and their capacity to lead to intelligent behaviour. As Fodor (1985, p. 93) once put it, the fundamental breakthrough of the Cognitive Revolution in comparison to past (representational) theories of mind was the adoption of the notion of computation to explain cognitive processing.

Naturalising physical computation is often taken to be crucial for the project of the cognitive sciences[1]. Given its central place in the computational-*cum*-representational basic framework of the cognitive sciences, philosophers aim to produce naturalistic theories that yield a robust, objective, non-trivial notion of computation in physical systems. Attacks against the feasibility of this project have been taken to cast doubt on the foundations of the cognitive sciences (e.g. Putnam 1988, Searle 1992, Van Gelder 1995)[2]. The underlying worry is that without such a naturalistic account of computation, mainstream cognitive science would be missing one of its sustaining pillars, the computational theory of cognition.

In the past decade, a particularly promising robust theory of computation has gained traction: the mechanistic view of computation (Piccinini 2015, Milkowski 2013, Fresco 2014, Coelho Mollo 2018). Piccinini's version of the view, in which appeal to natural teleology plays a crucial role, has been especially influential. The mechanistic view seems particularly well-equipped to satisfy three important *desiderata* for theories of physical computation, namely i) avoiding pancomputationalism — the claim that everything computes; ii) making space for the normativity of computation, i.e. the distinction between correct computation and miscomputation; and iii) adequately capturing scientific practice and the explanatory role computation plays in such practice.

Recently, an alternative approach to the problem of physical computation has been defended by Schweizer (2016, 2019) and Dewhurst (2018*a*): computational perspectivalism. According to computational perspectivalism,

[1]In what follows I often refer to physical computation simply as 'computation', as my concern in this paper is exclusively with computation in physical systems (rather than in mathematical theory).

[2]Sprevak (2018) offers a good critical overview of these attacks and of the most popular responses.

the nature of computational systems and processes hinges on explanatory perspectives that observers take toward physical systems. Dewhurst's and Schweizer's versions of computational perspectivalism are importantly different. One key difference is that while the former endorses a non-teleological version of the mechanistic account, the latter fully rejects the mechanistic view. Crucially, both invite a fundamentally different view of the nature and explanatory role of the notion of computation in science to that offered by teleological versions of the mechanistic account.

In this paper, I argue that computational perspectivalism fails to be a compelling alternative to objective (i.e. perspective-independent) views of the nature of computation, and especially to the teleo-based mechanistic view. I examine several versions of computational perspectivalism, showing that they either collapse into existing non-perspectival views; or end up with unsatisfactory or implausible accounts of physical computation. The teleo-based mechanistic view is therefore to be preferred.

Here is how I will proceed in what follows. In section §2 I briefly present the teleo-based mechanistic view of computation, which is the most plausible theory of computation on offer, and the framework to which computational perspectivalism aims to be a more compelling alternative. I then explore some varieties of perspectivalism in section §3, before examining computational perspectivalism and its virtues for what regards fulfilling the *desiderata* for theories of physical computation in section §4. Finally, in section §5 I argue that all varieties of computational perspectivalism have important shortcomings, making the approach less compelling than standard teleo-based mechanistic views.

2 The Mechanistic View of Computation

The mechanistic view of physical computation is an application to the notion of computation of the neo-mechanistic approach to science and scientific explanation, which has been very influential in the past few decades (Machamer et al. 2000). The basic idea behind the mechanistic framework is that (much of) science aims at revealing the mechanisms responsible for phenomena. Scientific explanation consists (largely) in identifying the relevant mechanism that explains a certain phenomenon of interest. This requires identifying the parts of the mechanism, what they do (i.e. their activities), and how they are organised. Mechanisms are always mechanisms for a certain phenomenon. Mechanistic decomposition reveals the componential and functional organisation of a system relevant for bringing about the phenomenon of interest[3].

Parts of mechanisms have systemic functions, i.e. their specific causal contributions to the behaviour of the overall mechanism. They may also have teleological, or 'proper' functions. That is, they may have functions that they are in a certain sense supposed to perform. Artefacts are a clear example of systems that have teleological functions, or teleofunctions for short. An electronic calculator is supposed to perform some mathematical operations on the inputs we provide. There is something it is for the calculator to work properly, i.e. when it correctly carries out the operations it is supposed to carry out, thus fulfilling its teleological function. And there is something it is for the calculator to work improperly, i.e. when it fails to carry out the appropriate operations, thus failing to fulfil its teleological function[4]. Even in the latter case, the calculator can be mechanistically decomposed, and its parts ascribed systemic functions that help explain how come it generates the (incorrect) outputs that it does. Teleological functions involve a sort of normativity; a claim about what a system is supposed to do, which is to some extent independent of its occurrent causal powers and the causal roles of its parts, i.e. their systemic functions.

While many or most mechanisms do not possess teleofunctions (e.g. weather systems), according to influential versions of the mechanistic view, some do (e.g. calculators, hearts). They are teleofunctional mechanisms (Garson 2013, Piccinini 2015). Piccinini's influential mechanistic account of computation has it that computational systems are teleofunctional mechanisms that possess the teleofunction to perform computations. It is a realist theory of computation, which upholds the objective existence of physical computation and of teleological functions. Although there is considerable variation among different versions of the teleo-based mechanistic view, the following individuation conditions seem largely to be shared among proponents (adapted from Coelho Mollo 2018, p. 3487).

Physical computation consists in:

[3]Though typically the mechanistic view is framed as being committed to realism about scientific posits, it need not be (Colombo et al. 2015).

[4]Artefactual teleofunctions, in contrast to the teleofunctions of biological systems, may partly depend on the intentions of designers and users. This does not make such teleofunctions less objective, since they hinge on facts of the matter about what intentions were at play in designing and/or using the system. Spelling out the specific role of intentions, if any, in determining such teleofunctions is an aim of theories of artefactual function (see Artiga 2016).

1. Manipulation of medium-independent vehicles according to a rule sensitive only to their degrees of freedom;
2. The medium-independent vehicles are components of a teleofunctional mechanism;
3. The manipulations that vehicles undergo are activities internal to a teleofunctional mechanism;
4. It is one of the teleological functions of the teleofunctional mechanism to carry out 1.

Two notions play a crucial role in the account: teleological functions, and medium-independence. A physical system must have the teleofunction to compute, if it is to count as computational (Coelho Mollo 2019). Moreover, its vehicles are individuated in medium-independent terms (Haugeland 1985)[5]. This means that the behaviour of computational systems is best captured by factors that do not depend on the details of the physical constitution of its parts. Rather, the relevant factors are the dimensions of variation of physical variables of the components — their degrees of freedom. Thus physical systems made of different materials can be computational; and may perform identical computations despite implementational differences provided that their physical dimensions of variation are relevantly similar[6].

2.1 Teleological Functions and Meeting the Desiderata

Teleology is a notion that calls for naturalisation, if it is to be legitimately employed in our scientific and philosophical theories. It involves ideas of purpose and normativity that seem *prima facie* difficult to explain in naturalistic, non-intentional terms. Naturalistic theories of teleological function try to dispel this impression. They can be divided into two broad families: dispositional theories, and selected-effects theories.

Dispositional theories bestow teleological functions on systems and their parts in light of their causal dispositions. Among dispositional theories,

[5]The notion of medium-independence is not uncontroversial, especially for what regards how stringent it is and which kinds of properties count as medium-independent. In this paper, I will take for granted that the notion is sufficiently stringent as to set it apart from other notions, such as weaker forms of multiple realisation. See Piccinini (2015, chap. 7) and Coelho Mollo (2019) for extended discussion.

[6]The latter point has been put in doubt by Haimovici (2013) and Dewhurst (2018*b*), a move criticised by Coelho Mollo (2018) and Fresco & Milkowski (forthcoming).

goal-based theories have received the most favour. According to them, the teleofunctions of a system and/or its parts are determined by their dispositions to contribute to a set of privileged goals of organisms — typically survival and inclusive fitness (Boorse 1976, Maley & Piccinini 2017).

Selected-effects theories, on the other hand, bestow teleofunctions on systems and their parts in light of the causal contributions of past instances of the same type of system and/or part, contributions which account for the selection and persistence across time of instances of the type (Millikan 1984, Neander 1991). While natural selection tends to be most often appealed to, several kinds of selection processes may be relevant. A promising recent view is Garson's (2011), which accords natural selection as well as other kinds of selection processes a role in the bestowal of teleofunctions to systems.

The teleo-based mechanistic view of computation need not be committed to a specific theory of teleological functions. It does however commit to there being a robust, objective theory of teleological functions able to ground the overall realist take on computation that it defends (Coelho Mollo 2019). This is a vulnerability of the view, given that its success hinges on whether a robust naturalistic theory of teleology can be provided. In section §5, I give defeasible reasons to lend credence to the latter claim, although I remain neutral on what naturalistic theory of teleofunctions is more likely to succeed.

Appeal to teleofunctions plays an important role in the mechanistic account, insofar as it helps to avoid pancomputationalism, i.e. the claim that all or most physical systems compute. Pancomputationalism risks making the notion of computation trivial, and its explanatory role vacuous: if everything computes, an explanation of how appeal to computation is nonetheless useful and informative is called for[7]. Moreover, pancomputationalism does not sit well with scientific practices in the computer sciences, which distinguish between systems that actually compute from those that can be merely computationally described or modelled[8]. While the latter category includes a large domain of physical systems, the former is taken to be much more restricted. Theories of physical computation that avoid pancomputationalism should thereby be preferred. According to the teleo-based mechanistic view, only a subset of mechanisms are teleofunctional. In turn, only a small subset of the latter are teleofunctional mechanisms that

[7]As we will see in section 4.1, limited forms of pancomputationalism avoid triviality.

[8]See Shagrir (1999) for an overview of earlier proposals on how to draw the distinction. Shagrir defends a semantic criterion, which is at odds with standard versions of the mechanistic view of computation.

possess the teleological function to compute. It follows that few systems are computational, whereby the mechanistic view succeeds in avoiding pancomputationalism, and does justice to the practices of the relevant sciences.

Physical computation plausibly involves a form of normativity. That is to say, for a system to be computational there must be something it is for it to compute correctly, and something it is for it to compute incorrectly, or miscompute. This is not an uncontroversial claim. Causal accounts, for instance, have it that physical computation captures abstractly described causal structure(s) that map onto a computational formalism (Chalmers 2011). There is no sense in which a causal structure, as causal structure, can succeed or fail. The computational rules that such systems follow are mere causal regularities mappable onto an abstract computational description at a certain point in time (Chalmers 2012, p. 230). In consequence, causal views tend to deny that making space for miscomputation is essential to an account of computation.

I take this to be ill-advised. The notion of computation is interesting to philosophers and scientists alike because it seems to have a distinctive explanatory role; a role that is not identical to that of other concepts that we already possess, such as that of abstract causal structure. Moreover, appeal to miscomputation plays important explanatory roles in computer as well as in cognitive science, as computational neuroscience and computational psychiatry make vivid (Fresco & Primiero 2013, Wang & Krystal 2014, Colombo forthcoming).

A satisfying account of physical computation — as *desideratum* (iii) prescribes — should vindicate the foundational and explanatory distinctiveness of the notion, making clear how it can play the roles it does as a theoretical posit in the computer and cognitive sciences. Accepting that computation involves a form of normativity, that computations can go wrong (and thus also right), seems essential to fulfil these requirements. It sets computation apart from causal structure, given the latter's lack of normativity, and accommodates the explanatory appeal to miscomputation in our sciences. Making space for miscomputation is thereby a plausible *desideratum* for theories of physical computation.

Computational processes respect rules that fix the transitions between inputs, internal states, and outputs. This means that a physical computational system may for a variety of reasons fail to respect such rules. Such failure may be of at least two kinds[9].

[9]For a much more sophisticated taxonomy of kinds of miscomputation, see Fresco &

It may be a failure of some component that makes it so that the system simply fails to compute — as for instance if the microprocessors of my electronic computer were to stop working (an example that I hope will always inhabit only the counterfactual realm). This is not a proper case of miscomputation, since no computation, and therefore no miscomputation, is being performed.

More interestingly, the failure may be such that the system computes, but fails to follow the computational rule it was supposed to follow. For example, my word processing software may, due to some internal defect, take my pressing of the ‘a’ key in the keyboard as input and generate as output in my display the symbol *s*, where the function it should have computed was from an ‘a’ key press to the displaying of the symbol *a*. For miscomputation to be possible, there must be a gap between what the system does and what it is supposed to do, such that these can match (correct computation) or mismatch (miscomputation)[10]. The foregoing mechanistic account of computation, given its reliance on teleological functions in the individuation of computational systems, naturally makes space for miscomputation. Whenever a computational system computes a function different from the one it is its teleofunction to compute, there is an instance of miscomputation.

3 Varieties of Perspectivalism

Dewhurst (2018*a*) has recently suggested that the mechanistic view of computation should not include a commitment to teleological functions in individuating computation. The notion of teleofunction, according to Dewhurst, has problematic naturalistic credentials, clashes with the notion of function most useful to cognitive and neuroscience — namely systemic functions (see section §2) — and at any rate represents an ontological commitment that the mechanistic view can do without — and of course the fewer problematic ontological commitments, *ceteris paribus*, the stronger the theory. Dewhurst recommends instead a perspectival approach to mechanisms and functions, following Craver (2013). Schweizer (2016, 2019) has argued for an alternative perspectival view of computation, which is not committed to the mechanistic account. My aim for the remainder of this paper is to show that perspectival-

Primiero (2013). For my purposes in this paper, the coarse distinction below suffices.

[10]See Tucker (forthcoming) for an extended argument for the need of such a gap, and more generally for a compelling and detailed account of miscomputation within the mechanistic view.

ism is not a promising approach to computation. Before moving to that, we must be clear on what computational perspectivalism, in its many varieties, involves.

Perspectivalism (or perspectivism) in philosophy of science encompasses a variety of positions that go from the rather trivial to the ontologically revisionary, with Kantian-flavoured views somewhere in between[11]. What they share is the idea that knowledge of the world is mediated, in more or less essential ways, by theoretical and pretheoretical assumptions, mathematical and experimental apparatus, explanatory interests and purposes, which come to compose a perspective that individuals and/or scientific communities take toward the world. That thesis, by itself, is uncontroversial. There is little doubt that our scientific knowledge of the world and that our methods for expanding it involve taking perspectives so understood. As philosophers of science have insisted for quite some time, there are no unmediated ways of deriving theories from 'raw' data. Epistemic perspectivalism of this type is innocuous, and widely accepted.

There are less innocuous versions of epistemic perspectivalism. Stronger forms of epistemic perspectivalism take the mediation of our knowledge of the world by perspectives to place important constraints on what we can know. By their lights, some or all of our scientific knowledge is intrinsically perspectival, and we are barred from gaining knowledge about perspective-independent properties of the world. In other words, in some or all domains of science, the knowledge we get can never reach, as it were, 'how things really are', being trapped within 'how things look like from here'.

A yet stronger type of perspectivalism is ontological. The rough idea is that some or all facts are perspective-dependent; not only our knowledge of them. Ontic perspectivalism entails non-innocuous epistemic perspectivalism: if the nature of entities and properties in the world is perspective-dependent, then our knowledge of them can only be perspective-dependent (Chakravartty 2010). Epistemic and ontic perspectivalism can be limited or universal: they can be true of some domains of knowledge or some parts of the world, while being false of other domains or parts. Given my purposes in the foregoing, I will focus only on perspectivalism about physical computation.

A further distinction that will play an important role in what follows is that between what I call 'thin' and 'thick' conceptions of explanatory perspective. Thin explanatory perspectives are relatively unconstrained.

[11]The taxonomy I offer here follows Chakravartty (2010). See also Giere (2006), Morrison (2011), Massimi (2018), Ismael (2016).

They are individuated simply by appeal to the interests and purposes of one or a few individuals. I can take the perspective of seeing my lamp as computing the identity function, because, say, I am interested in illustrating a distinction. Thin perspectives are highly arbitrary and very liberal: with the exception of basic constraints of individual rationality and responsiveness to evidence, there are few if any restrictions on the perspectives that may be taken toward (parts of) the world.

By contrast, thick explanatory perspectives are much more constrained, insofar as they embody richer epistemic structures[12], i.e. theoretical commitments, methodological and evaluative standards of (scientific) communities, modelling practices, explanatory aims, and so on (Massimi 2018). Thick perspectives are much more demanding and much less arbitrary than thin perspectives, inasmuch as they depend not only on the purposes and interests of single individuals, but rather on a sophisticated, socially created and maintained collective endeavour with stringent normative standards.

With these distinctions in hand, I will assess the plausibility of several varieties of computational perspectivalism, especially in light of the *desiderata* mentioned above, namely avoiding pancomputationalism, making space for miscomputation, and capturing the explanatory practices of the computer and cognitive sciences. I will argue that thick varieties succeed in respecting the first two *desiderata*, though as I will show in section §5, fundamental problems lurk elsewhere, especially when it comes to satisfying the third one.

4 Computational Perspectivalism

Dewhurst's (2018*a*) computational perspectivalism draws inspiration from a type of perspectivalism recently defended by Craver (2013), and which is motivated and finds its application within the neo-mechanistic approach to scientific knowledge and explanation. I will label it 'mechanistic perspectivalism'.

According to mechanistic perspectivalism, functions are not reducible to the causal structure of the world. They are rather ways of carving that complex and rich causal structure into mechanisms, in accordance with the instrumental and explanatory interests of agents. Mechanisms, as we have seen, are explanatory of capacities that systems display, of functions they

[12]From now on, I will often drop the qualifier 'explanatory' when talking about explanatory perspectives.

perform. Ascribing functions to parts of the world leads to carving its causal structure in terms of the entities and activities responsible for the capacities of interest. That is to say, ascribing functions filters the causal nexus so as to yield a mechanistic nexus: a mechanistic hierarchy with a specific capacity as its topping off point (Craver 2013). Function ascription depends on a perspective, on a choice about what features of the causal structure of the world to privilege. We take such perspectives in light of our explanatory interests and instrumental needs in each case.

Different perspectives lead to ascription of different capacities, and hence to the individuation of different mechanisms in the world that explain the performance of those capacities. Importantly, the structure of the world constrains what perspectives can be taken, what functions can be ascribed, and which mechanisms are relevant to explaining those functions. Given a function ascription, only one (or a few) mechanistic explanations will be acceptable. This is so because, according to Craver, given a certain *explanandum* only some ways of carving the perspective-independent causal structure of the world yield appropriate explanations.

Craver's mechanistic perspectivalism is a form of innocuous perspectivalism, inasmuch as it holds that the parts and activities appealed to in mechanistic explanations are perspective-independent features of the causal structure of the world. It is the way we group them together into mechanisms that is perspective-dependent, insofar as ascription of function is dependent on our explanatory interests and purposes. For instance, if I am interested in explaining the computational capacities of my laptop computer, I will carve it in terms of transistors, memory components, etc. If instead I am interested in explaining its capacity to disperse heat in the environment, I will carve it in terms of power sources, heat-generating components, fans, etc. While these two explanations individuate components within my laptop computer differently, they both capture aspects of its causal structure.

Teleological functions, on the other hand, do not capture anything objective in the world by the lights of mechanistic perspectivalism: "the causal structure of the world does not ground talk of goals, purposes, and preferential states" (Craver 2013, p. 147). Ascription of teleology to physical systems is something we do for pragmatic or heuristic reasons. There is no perspective-independent notion of teleofunction to be had, and thereby no purely objective grounds for claims about proper function and malfunction. Mechanistic perspectivalism endorses at most a sort of instrumentalism about teleofunctions, these being ultimately grounded on our choices, rather than on the objective organisation of the world.

Dewhurst (2018*a*) shares Craver's scepticism about teleological functions. For a sceptic about teleofunctions the teleo-based account of mechanistic computation offered in section §2 is wrong-headed. It relies on an objective, robust notion of teleological function for computational individuation. If there are no objective teleofunctions; if these are essentially interest-based heuristics for carving up causal structures, we cannot use them to ground an objective notion of physical computation.

Dewhurst suggests that the appeal to teleological functions be replaced with a perspectival notion of function. In his perspectivalist view, a mechanism is computational only relative to an explanatory perspective: a mechanism has the function to compute only "in contexts where this function contributes to our explanation of some phenomenon" (Dewhurst 2018*a*, p. 573). Computational systems are thus all those physical systems that possess the appropriate physical structure such that they are interpretable as performing computations from some explanatory perspective (Dewhurst 2018*a*, p. 581).

Let us see how this picture fares when confronted with the *desiderata* on theories of computation presented above.

4.1 Computational Perspectivalism and Pancomputationalism

Dewhurst seems to accept a rather liberal notion of explanatory perspective. For he claims that we can take a computational explanatory perspective toward (almost) any physical system, whereby a limited form of pancomputationalism follows from perspectivalism (Dewhurst 2018*a*, p. 585). There are always ways of carving a (minimally causally complex) physical system such that its states and processes can be seen as implementing computations (Chalmers 2011). It would follow that (almost) every physical system computes at least one function; or more weakly, houses at least one computational mechanism formed by a subset of its physical states. Moreover, a single system may be appropriately seen as performing more than one computation at the same time. Physical systems can be described at different causal levels, and each of these levels may be interpreted as performing one or more computations.

Nevertheless, there is a limited range of acceptable computational descriptions. The physical structure of the system constrains what computational descriptions are legitimate. For instance, computational descriptions that posit computational elements, such as digits, that do not match the physical structure of the system — e.g. because the system does not respond systematically to those elements, and thus does not follow a rule defined over them — are excluded. The constraints on legitimate computational explanations imposed by the physical structure of the system allow the resulting view of computational individuation to avoid triviality. Physical systems cannot be seen as performing every computation, given that most computational ascriptions would fail to respect the physical organisation of the system, or respect it to a lower degree than other, better ascriptions[13]. The account does not thereby fall prey to the more troubling, trivialising unlimited form of pancomputationalism, i.e. the claim that every physical system computes every function. Dewhurst (2018*a*) accepts these consequences of perspectivalism about computation, arguing that only unlimited pancomputationalism leads to triviality and jeopardises the explanatory purchase of computation. A theory of physical computation, he holds, can comfortably live with limited pancomputationalism.

Although I agree that limited pancomputationalism does not trivialise computation, I believe that computational perspectivalism need not lead to pancomputationalism, not even of the limited sort; thus sidestepping any worries about limited pancomputationalism failing to capture scientific practice or endangering the explanatory role of computation in computer and cognitive science. Limited pancomputationalism only follows from a thin view of explanatory perspectives on which they are determined by little more than some individual or other ascribing computational function to a system. In consequence, there are very few constraints on which systems can be seen computationally, and on what computational ascriptions are appropriate. However, this view of explanatory perspectives is not the kind of perspectivalism relevant to science. Explanatory perspectives are grounded on rich collective and historically-shaped scientific endeavours, with their specific theoretical assumptions, methodological lore, theoretical posits, and set of phenomena that are their explanatory targets. Explanatory perspectives, on this 'thicker' understanding, place considerable constraints on what features of the world can fall under their purview.

Perspectivalism in philosophy of science is typically concerned with thick perspectives. In consequence, I will focus on the stronger version of compu-

[13]Dewhurst (2018*a*) seems to prefer a criterion based on Kolmogorov complexity for measuring the degree of fit between computational ascriptions and the physical organisation of systems (see also Milkowski 2013). Assessment of whether this criterion or alternative ones are plausible is beyond the scope of this paper. I will thus assume throughout that some such criterion is adequate.

tational perspectivalism that follows from accepting thick perspectivalism, which, as I will show, has better prospects for meeting the *desiderata* of a theory of physical computation. This thicker view of explanatory perspectives invites the addition of further constraints that systems liable to be seen under the computational perspective need to satisfy. Under thick computational perspectivalism, computational systems are, as before, those systems to which we can take a computational perspective. But by enriching computational perspectives, not every system may be such that we can take a computational perspective toward it, hence avoiding limited pancomputationalism.

One first criterion for enriching the computational explanatory perspective is the appeal to medium-independence, which is part and parcel of the mechanistic view. The computational perspective applies to those capacities of systems for which an explanation in terms of medium-independent vehicles can be given. Digestive systems as digestive systems, for instance, are excluded. To explain digestion, one must arguably appeal to properties that can only be individuated medium-dependently, i.e. kinds of proteins, enzymes, etc. (Piccinini 2015, p. 146). Physical computation, in contrast, does not require specific kinds of physical properties to be explained. Although in each instance computational functions must be performed by a physical substratum, what that substratum is — silicon chips, neurons, vacuum tubes — is irrelevant to its being the performance of a computation. In brief, ascribing a computational function to a system contributes to explaining only those phenomena that do not call for medium-dependent mechanistic individuation of its parts and activities.

The inclusion of medium-independent individuation as a component of computational explanatory perspectives is a welcome addition to computational perspectivalism, but it does not help much with the problem of pancomputationalism. While digestive systems, *as digestive systems*, are clearly outside the purview of the computational perspective; it is much less clear that digestive systems, *as physical systems*, cannot be seen as implementing medium-independently characterised operations over vehicles. It seems possible to describe the causal structure of digestive systems in such an abstract way that no commitment to specific physical realisers is made (Chalmers 2011). If so, then digestive systems fall under the purview of computational explanation; although not as *digestive* systems, since the phenomenon that ascription of computational function will help to explain is not digestion, given the latter's medium-dependent nature. Rather, computational explanation will reveal the fine-grained input-output functions that the system implements when its causal structure is medium-independently described. Pure appeal to medium-independence, in sum, does not by itself avoid limited pancomputationalism.

These considerations point toward further constraints on computational perspectives. Despite it being arguably possible to explain the goings-on in digestive systems computationally, it does not seem that this explanation would be particularly useful or informative about their digestive capacities. Seeing digestive systems as computational does not seem helpful either instrumentally or theoretically: it does not correspond to the capacities that biology and physiology ascribe to such systems, and it does not allow us to intervene in useful ways on digestive systems. This intuition can be fleshed out by adding further requirements, such that the computational explanatory perspective only applies to systems to which it contributes explanations that respect plausible criteria of theoretical and instrumental adequacy.

Schweizer (2014, 2016) puts forward a list of factors that distinguish instrumentally and theoretically interesting ascriptions of computational nature from useless ones. Here, I will consider such factors to be embedded in the computational explanatory perspective, helping to individuate the systems that are adequately captured from that perspective.

First, computational ascriptions must lead to epistemic gain — they must be such that applying them to physical systems adds something relevant to our knowledge. This would rule out digestive systems, since we arguably do not gain new information about their workings by seeing them under a computational perspective, rather than under a non-computational one. Identifying their abstract causal structure contributes little if anything to our understanding of their relevant capacities in addition to what is contributed by a non-computational, medium-dependent approach; and neither does it augment, comparatively, our capacity to predict the future states of the system.

Second, systems helpfully seen under the computational explanatory perspective should also be relatively versatile and reliable. Computational explanation gets much of its purchase from its capacity to explain how physical systems can reliably produce specific outputs when fed certain inputs. There is little if any explanatory purchase in seeing my lamp as computing the identity function — it is not computationally versatile, though it be fairly reliable. By contrast, explaining the behaviour of my laptop computer in terms of its execution of a word processor has considerably more explanatory power — it helps answering the question of why the computer behaves as it does given the inputs I produce. In consequence, the latter, but not the former, is liable to being seen as performing computations.

Third, the computational explanatory perspective should also embody a manipulability constraint, insofar as the possibility of intervention and control helps to justify the explanatory and instrumental role of appeal to computation. The computational perspective can be taken only toward physical systems whose computationally individuated states can be intervened upon such that the computational function(s) they perform change reliably and predictably. While this constraint is plausible for what regards designed computers — after all, we build, modify, and repair them by means of such interventions — it may seem problematic when applied to cognitive systems, given our current ignorance of what their computational states and processes are. This worry can be accommodated by noting that the foregoing requirement is best seen as an 'in principle' claim. If the computational theory of cognition is correct, then a future, fully developed and instrumentally powerful cognitive science will be able to identify and intervene on those states and processes that carry out computations (*qua* computational states and processes), for instance for therapeutic reasons or for building artificial cognitive systems. If the computational theory of cognition proves to be false, then computational perspectivalism would reject the claim that the computational explanatory perspective should be taken toward cognitive systems (while it may preserve a role for computational *modelling* as a heuristic tool). These are the appropriate results, if our theory is to capture the distinction between computational and non-computational physical systems; and be neutral on empirical questions such as the truth or falsity of the computational theory of cognition.

On the foregoing thick version of computational perspectivalism, we can agree with Dewhurst that to be a computational system is to be ascribed from an explanatory perspective a computational function that helps explain a certain phenomenon. But we should resist the claim that it follows that every physical system can be examined under that perspective. Rather, few systems are fit for explanation under the computational perspective, once it is appropriately enriched. Limited pancomputationalism thereby does not follow from (thick) computational perspectivalism.

4.2 Computational Perspectivalism and Miscomputation

Does the thick version of computational perspectivalism that I have been examining make space for miscomputation?

I think it does. Factors such as the explanatory targets and the accepted conceptual and methodological framework of a certain domain of science, as we have seen, partly constitute a (thick) explanatory perspective. Such factors, together, ground a notion of what a certain type of system is supposed to do. For under a certain explanatory perspective, a type of system will be seen as contributing or failing to contribute to the phenomenon or capacity that is the target of explanation, while the overall research framework will help determine the proper way of individuating the system, its causal interactions and roles, and its proper decomposition into parts and their activities[14].

Take computational neuroscience. At the basis of most of the efforts in the field lies the idea that brains are computational systems (Kriegeskorte & Douglas 2018). That foundational hypothesis is based, at least largely, on independent theoretical considerations as well as on inference to the best explanation. The notion of computation plays a central role in explaining the capacity of cognitive processes to generate appropriate behaviour on the basis of external inputs (e.g. sensory information) and internal states (e.g. beliefs about the world), an achievement that no other theoretical posit has to date been able to accomplish (Carandini & Heeger 2012, Wang & Krystal 2014).

Given the explanatory target of computational neuroscience, i.e. explaining by means of computational states and processes how organisms behave, the specific research projects it pursues have as their starting point a certain cognitive capacity that organisms display, and which is to be, at least partly, computationally explained. For instance, given the ability of certain kinds of organism to acquire information about the shape of objects from visual stimuli, it can be hypothesised that part of the explanation is that mechanisms within the visual system have the function to perform computations that extract contour information from shadow information. Experimental and modelling work can confirm or disconfirm the hypothesis, as well as help individuate the mechanisms that perform the relevant computations (e.g. Lehky & Sejnowski 1988, Khuu et al. 2016). The scientific framework also has the tools to type mechanisms and their functions in such a way that they can be generalised to a more or less extended domain of types of organism. Contributions from other disciplines, such as evolutionary biology and ethology, may suggest possible selection processes or contributions to objective goals that lend force to ascribing teleofunctions to the relevant

[14] Thus understood, Craver's (2013) mechanistic perspectivalism is thick.

mechanisms. Therefore, if future work should establish that visual systems include mechanisms whose function it is to perform those computations, organisms belonging to the type whose relevant mechanisms do not extract contour information from shadow information — but rather compute some other function — miscompute.

In brief, thick computational perspectivalism has the tools both to determine, in light of conceptual and methodological considerations, what systems are to be seen as computational under a non-arbitrary explanatory perspective; as well as which function ascriptions are appropriate given the explanatory aims and classificatory practices embedded in the explanatory perspective.

It is possible that different (thick) explanatory perspectives may ascribe different computational functions to the same mechanisms, perhaps even in such a way that they take as instances of correct computation what current computational neuroscience sees as instances of miscomputation. Similarly, within computational neuroscience incompatible hypotheses about the computational functions of certain brain mechanisms may be put forward. From this, however, it does not follow that the notion of miscomputation is trivial, or arbitrary. For any alternative framework would be in direct competition with current computational neuroscience. Its different methodological and classificatory practices, which lead to the alternative function ascriptions, would need to be assessed and compared to the current framework in light of the preferred theoretical, methodological, and evidential virtues that the relevant scientific communities adopt as criteria for theory choice. Analogously, incompatible hypotheses within computational neuroscience would be in competition, and the choice between them dictated by the criteria for theory choice of the field. What would trivialise miscomputation is the claim that there is no non-arbitrary way of distinguishing appropriate from inappropriate computational function ascriptions. Showing that there is competition between scientific frameworks or hypotheses falls considerably short from lending force to that claim (Morrison 2011)[15].

In sum, thick computational perspectivalism can make space for miscom-

[15] As an aside, note that thin computational perspectivalism has trouble giving a non-trivial account of miscomputation. For given the few constraints imposed on perspectives and what falls under them, one is free to prescribe different computational functions to the same system under the same mechanistic decomposition — i.e. preserving all individuated parts and activities. It follows that any physical system can legitimately be seen as correctly computing a function f, and equally legitimately as miscomputing an indefinite number of functions g, h, q …

putation, thereby fulfilling the relevant *desideratum* on theories of physical computation. What remains to be seen is to what extent computational perspectivalism differs from standard, teleo-based theories of computation. And if it does differ, what species of perspectivalism it is: innocuous, epistemic, or ontic. It is by considering these questions, I suggest, that the problems with computational perspectivalism come to surface. Even though it has the tools to satisfy *desiderata* (i) and (ii), it has trouble with descriptive accuracy and explanatory adequacy as prescribed by (iii).

5 The Varieties of Computational Perspectivalism Assessed

Computational perspectivalism has it that attribution of computational functions, and thereby of computational nature to physical systems depends on taking a specific kind of explanatory perspective toward them. As we have seen, a thick version of the view is particularly attractive, since it seems to do away with the potentially problematic appeal to objective teleofunctions, whilst keeping at bay pancomputationalism and making space for miscomputation. In this section, I examine whether these apparent advantages survive a closer look. I argue that computational perspectivalism ends up in an uncomfortable dilemma. If it tries to be an innocuous kind of perspectivalism, it either collapses into causal views of computation, or into the teleo-based mechanistic account. If on the other hand stronger forms of perspectivalism are upheld, the resulting views suffer from important shortcomings.

The central criterion that distinguishes innocuous from non-innocuous types of perspectivalism is the commitment, or lack thereof, to there being perspective-independent facts about the world, and about which we can acquire knowledge. In the case of computational perspectivalism, the claim is limited to computational facts. Innocuous computational perspectivalism has it that by means of the relevant thick explanatory perspective(s) we acquire knowledge about the perspective-independent nature and structure of computational physical systems. Epistemic computational perspectivalism is neutral about the perspective-independent existence of computational physical systems, for it holds only that we cannot know any perspective-independent facts about them. Ontic computational perspectivalism, finally, claims that the notion of computation is ontologically dependent on the perspectives that we take, and thereby that computational systems, as computational, have no perspective-independent existence.

Let us now turn to the relations between the varieties of computational perspectivalism and the other accounts currently on offer, with an eye to assessing their relative promise.

5.1 Innocuous Computational Perspectivalism

According to innocuous computational perspectivalism, computational explanations (i.e. explanations from a computational explanatory perspective) tap into objective, non-perspectival features of the world. First, they capture, as a minimum, perspective-independent aspects of the causal organisation of physical systems: the causal regularities instantiated by types of internal states and their transitions describable by a computational function. Second, as we have seen in section 4.2, thick explanatory perspectives also include prescriptions about how a computational system should behave, i.e. what its proper or normal behaviour is. This introduces a sort of normativity into the notion of computation, making space for miscomputation. Varieties of innocuous computational perspectivalism all agree on the former claim, while they differ for what regards the latter. More precisely, there are two stances that an innocuous perspectivalist may take toward the normativity of physical computation.

First, a computational perspectivalist may deny that normativity is essential to the notion of computation, seeing function ascription as merely heuristic. This seems to be the view that Dewhurst (2018*a*) prefers. As I argued in section 2.1, however, this position is unsatisfactory. It is equivalent to a purely causal theory of computation, as the one defended by Chalmers (2011), whereby computation is little more than a label for abstract causal structure. As we have seen, on causal theories the notion of computation has no distinctive explanatory value over and above that of abstract causal structure. Moreover, denying that computation is normative fails to make space for miscomputation, and thus does not do justice to our best computer and cognitive sciences, in which miscomputation plays an important explanatory role. For these reasons, I take that this form of innocuous perspectivalism, which is actually a causal theory of computation, should be rejected.

A second option is to hold that normativity is essential to physical computation, and that function ascriptions made from the relevant explanatory perspective(s) capture perspective-independent features of computational systems. This view has much to recommend it, but it is indiscernible from the mechanistic account presented in section §2. For it accepts that the relevant explanatory perspectives reveal both the perspective-independent causal structure of computational systems, which mechanistic decomposition captures; and the teleofunctions to compute that those systems have independent of our explanatory perspectives, which are captured by theories of teleological function. This second sort of innocuous computational perspectivalism, as much as the first, is not a distinct theory. It is just the kind of innocuous perspectivalism that the teleo-based mechanistic view accepts (Piccinini 2015, p. 142).

Another option is to accept that normativity is essential to the notion of physical computation, but hold that it depends intrinsically on explanatory perspectives. This kind of perspectivalism is no longer innocuous. For if it claims that computation essentially involves miscomputation, and that miscomputation essentially depends on explanatory perspectives, then the notion of computation may not capture perspective-independent features of the world. A strong form of perspectivalism is thus at play, to be examined in the next section.

In sum, innocuous computational perspectivalism, *contra* Dewhurst (2018*a*), is not really in the cards. It collapses into existing theories of computation, either causal or teleo-mechanistic. The perspectivalist about computation has other options available, but they are far from innocuous.

5.2 Non-Innocuous Computational Perspectivalism

Epistemic computational perspectivalism is the view that we cannot acquire perspective-independent knowledge about physical computational systems. The view only makes an epistemic claim, and is therefore neutral on the ontological status of physical computation. There are two main grounds for holding epistemic computational perspectivalism: a) because the notion of computation is instrumentally or pragmatically useful, but may not correspond to actual features of the world; or b) because computation is ontologically dependent on explanatory perspectives. The former is identical, or close enough, to instrumentalism about computation. The latter is identical to ontic perspectivalism. In what follows I will examine these two varieties of non-innocuous computational perspectivalism in turn.

Computational Instrumentalism

Schweizer (2016, 2019) advocates an explicitly instrumentalist theory of computation. According to him, the notion of computation does not capture objective features of the world, but is nonetheless a useful conceptual tool that plays a heuristic role in several scientific fields — especially in the computer and cognitive sciences. Performing computations is not a property that physical systems have objectively, but rather depends on an 'purely observer-dependent act of ascription'. Such computational ascriptions are constrained in their scientific appropriateness by pragmatic constraints that make appeal to computation more or less useful — the most central of which were introduced above, in section §4. For Schweizer, application of some or all of these constraints depends on the explanatory context. Different explanatory and instrumental contexts motivate the employment of pragmatic constraints which may play no role in other contexts.

Even in those contexts in which application of the notion of computation is explanatorily useful, there is no ontological commitment to the systems under investigation performing computations. Although these considerations are compatible with a neutral position on the (perspective-independent) existence of physical computations, Schweizer rejects the idea that there are physical systems that actually compute. His instrumentalism is underlain by ontological eliminativism about computation. His main argument for this claim is a familiar one. He agrees that miscomputation is essential to computation, that is, that the notion of computation is cogent only if it includes a normative element. However, like Craver (2013), Schweizer is sceptical about natural teleology. Therefore, the normativity of computation, according to him, can only come from prescriptive interpretations that observers impose on physical systems. Physical computation cannot thereby be a perspective-independent feature of systems in the world.

There are several problems with instrumentalism about computation. First, one of its underlying motivations, i.e. scepticism about natural teleology, seems problematic when we look at the explanatory practices of the life sciences. The notion of teleological function is part and parcel of several of our mature biological sciences, as is the appeal to normative considerations to explain deviations from appropriate functioning in terms of abnormal internal workings or abnormality in the environment. As we have seen, there are several plausible proposals on how to naturalise teleology. Schweizer (2019) quickly dismisses these attempts, in a way similar to Craver (2013), on the grounds that teleofunctions are underlain purely by causal-historical factors. However, this is hardly a counterargument: the project of naturalisation is exactly to rid the notion of 'anthropomorphic heuristics', grounding it in causal features of the world. Borrowing Jerry Fodor's famous phrase (originally about naturalising intentionality), if teleofunctions are real, they must be really something else.

For a sceptic about teleofunctions, the considerations above may be unconvincing. They may think that the promise of current theories of natural teleology is only apparent and will not hold out. This sort of sceptical challenge is notably difficult to counter a priori. Although I cannot offer a full defence of natural teleology here, the fruitfulness of the notion in scientific practice, as well as the philosophical consensus around its explanatory role (if not its nature), provide at least *prima facie* reasons to place the burden of proof on the sceptic[16].

Second, analogous considerations from explanatory adequacy and descriptive accuracy regarding scientific practice suggest another way in which perspectivalism is problematic. The computer and cognitive sciences seem, at least often enough, to be committed to the claim that the systems they build and investigate are really computational; in contrast to systems that can merely be modelled computationally. By denying that computations are objective features of the world, instrumentalism about computation is at odds with these practices. If we consider explanatory adequacy and descriptive accuracy to be relevant criteria for philosophical theory choice, as per *desideratum* (iii), realist views of computation such as the teleo-based mechanistic account have the upper hand[17].

Instrumentalists about computation can reply in two ways. First, they may endorse a general instrumentalist view of science in which the theoretical posits of our best sciences do not involve ontological commitment, making the considerations above void. Although this is an acceptable move, it makes *computational* instrumentalism lose much interest, given that it simply flows from a much broader instrumentalism about scientific posits — assessment of which is well beyond the scope of this paper. Second, and more interestingly, computational instrumentalism may be defended by showing: i) that there are good reasons for withholding ontological commitment to computation specifically; and ii) that other factors make it so that the notion is nonetheless ineliminable from our best explanations in computer and cognitive science[18]. Let us start with ii), which I believe helps shed light on i).

[16] For analyses and defence of the role of teleofunctions in the life sciences as well as of influential naturalistic theories of teleology, see Godfrey-Smith (1993), Griffiths (1993), Huneman (2013), Neander (2017).

[17] Sprevak (2018) also appeals to explanatory adequacy and descriptive accuracy in criticising anti-realist approaches to computation in the cognitive sciences.

[18] I am indebted to an anonymous referee to this journal for pressing me on this point.

As we have seen, Schweizer claims that appeal to computation is useful insofar as it allows better prediction, intervention, and control when applied to complex physical systems that feature certain types of flexible and reliable behaviour. Applying the notion of computation to such systems leads to epistemic and instrumental gains, justifying the presence of the notion in scientific explanations despite the denial of ontological commitment to computation. Thus, the instrumentalist may reply, we need not ontologically commit to computations in order to recognise the explanatory role and purchase of the notion in our best sciences.

I take this line of reasoning to be unconvincing. The epistemic and instrumental advantages that computation provide us with are themselves in need of explanation. If those advantages are due to the fact that the notion best captures actual features of systems relevant for some kinds of explanation, e.g. how they reliably and flexibly transition from a variety of different inputs to task-appropriate outputs, then it becomes unclear in which sense ontological commitment to computation is being denied. It may turn out that such epistemic and instrumental considerations lead to different views of the nature of physical computation than the ones currently on offer, but the resulting view will nonetheless be realist.

It may be replied that the instrumentalist actually makes a weaker claim than the one I suggest above — a claim about usefulness, while I helped myself to the idea that computation *best* captures some features of the world. But this reply is of no help, for it does not explain why our best explanations in computer and cognitive science appeal to physical computation. Appeal to mere usefulness without ontological commitment implies — short of general scientific instrumentalism — that the computer and cognitive sciences can (and should) come up with better explanations that involve only posits that we ontologically commit to. The instrumentalist about computation, however, has provided no clear guidance about what those explanations should look like.

Moreover, the principle of inference to the best explanation (IBE) spells trouble for instrumentalism: if positing computations is the best way to explain some phenomena, as least at the current state of our best sciences, then we should, as a working hypothesis at least, commit to their existence. The computational perspectivalist can resist this argument only by either denying the principle of inference to the best explanation in general, thus sliding toward a general instrumentalist view of science as a whole; or denying that IBE applies to the notion of computation in particular.

The former strategy, as we have seen, is not particularly interesting for a theory of physical computation. The latter, on the other hand, is motivated by two claims that we are already familiar with, and that bring us back to i) above: computation is normative, and there are no sources of normativity beside perspectival ascription. Therefore, theoretical posits that appeal to normativity, such as computation, are non-innocuously perspectival. Scepticism about natural teleology, as I pointed out, is at the heart of computational perspectivalism. Not much else speaks in favour of the view, as I hope to have shown. But I hope to have also shown that this is a rather problematic foundation for non-innocuous computational perspectivalism. Although the philosophical controversy about natural teleology is far from solved, there are at present good, though defeasible reasons for accepting the claim that teleofunctions are *bona fide* features of the world.

Even if one should prefer to remain neutral about natural teleology, attention to practice in the relevant sciences, and especially to the role of IBE in them, suggests that we should ontologically commit, as a working hypothesis, to physical computation. As Sprevak (2013) points out, most of the knowledge acquired by the cognitive sciences stems from accepting the posits that best explain the relatively partial and indirect experimental data available. Computation is one of these posits. There is not much left to guide cognitive science and justify its practices and results if IBE is (even partially) abandoned.

Finally, computational instrumentalism, if wedded to ontological eliminativism, a priori forbids a computational theory of cognition, a consequence that both Dewhurst (2018*a*) and Schweizer (2019) endorse. For if computation is only in the eye of the beholder, it cannot be at the foundation of what allows observers to behold, interpret, and ascribe — on pain of vicious circularity. At best, we can see the computational view of cognition as a useful fictional model, but not as a hypothesis about the nature of cognitive systems and of how they bring about intelligent behaviour. The hypothesis that cognition and intelligence are made possible partly by the performance of computations, one of the bases of cognitive science (and of much philosophy of mind), is taken out of the table by pure armchair reflection. This is a problematic way of rejecting an empirical hypothesis, oblivious to its scientific fruitfulness, (dis)confirming evidence, and explanatory purchase[19].

In brief, instrumentalism about computation is problematic, especially when compared to ontologically committed accounts of computation, such

[19]Ontologically neutral computational instrumentalism, on the other hand, leads to permanent quietism about the computational theory of cognition, a similarly problematic stance to take toward an empirical hypothesis.

as the teleo-based mechanistic view. The latter, in contrast to the former, is descriptively adequate, capturing the usage of the notion of computation in the computer and cognitive sciences; it is explanatorily adequate, doing justice to the explanatory practices and commitments of our best sciences; and it does not preclude a priori the truth of the driving hypothesis of cognitive science, i.e. that cognition is at least partly computational. These arguments are perhaps not as decisive as a realist may have wanted, but I think they are enough to shift the burden of proof to the computational instrumentalist, and to justify lending more credence to realism about computation than to instrumentalist views.

Ontic perspectivalism

The other non-innocuous option for the computational perspectivalist fares no better, and for largely similar reasons, so my treatment here will be brief.

On ontic perspectivalism about computation, facts about computational systems and processes are intrinsically dependent on perspectives, there being no non-perspectival facts about physical computation. Although ontic perspectivalism, in contrast to instrumentalism, ontologically commits to computation, it is unclear that it can be made plausible. Ontic perspectivalism claims that physical computation exists, but that its existence depends on perspectives we take toward the world. However, physical computation, in contrast to money, universities, social and political institutions, seems to be the wrong kind of thing to have its existence depend on mental states and/or social practices of sentient beings, or on the scientific perspectives they take[20]. In a weaker sense, the notion of computation does depend on our explanatory practices, on scientific communities, and so forth. The concept has been developed within certain explanatory projects, and its application in various branches of science depends on standards of scientific and explanatory adequacy. However, our sciences are mostly committed to the claim that natural and artificial computational systems exist, and would exist even in the absence of individuals and communities taking perspectives toward them.

[20]If the only computational systems in the universe are cognitive systems and artificial computers designed by them, then there is a kind of ontological dependence at play: computational systems would not have existed were it not for the existence of cognitive systems. However, this is not the relevant kind of dependence. Ontic perspectivalism claims that computation is ontologically dependent on explanatory perspectives taken by cognitive systems, not on the existence of cognitive systems per se.

Moreover, even if ontic perspectivalism about computation could be made plausible, it would share most of the shortcomings of computational instrumentalism. It precludes the hypothesis that cognition is computational, on pain of circularity; and it does not square well with the ontological commitments of the computer and cognitive sciences, since they do not typically take the existence of computational systems to depend on explanatory perspectives we take.

In sum, non-innocuous forms of computational perspectivalism are unsatisfactory. They fare badly when compared to teleo-based mechanistic views of computation. Eschewing natural teleology from the account of physical computation proves to be too high a price to pay for little gain.

6 Concluding Remarks

In this paper, I have argued that perspectival theories of computational individuation, as currently proposed, are not satisfactory. Innocuous computational perspectivalism is indiscernible from existing realist views of computation, of which the most plausible is the teleo-based mechanistic account. Non-innocuous versions of computational perspectivalism, on the other hand, have several important shortcomings. The mechanistic view is therefore to be preferred, despite its richer ontological commitments.

Acknowledgments

I am indebted to Matteo Colombo, Joe Dewhurst, Margherita Arcangeli, three anonymous referees to this journal, as well as to audiences at the Computation and Causation in Cognitive Neuroscience Workshop 2018, in Cologne, and at the Philosophy of Mind colloquium of the Berlin School of Mind and Brain for helpful comments on previous versions of this material. This research was supported by the Berlin Center for Advanced Neuroimaging, Charité-Universitätsmedizin Berlin, and by the Science of Intelligence Cluster (funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2002 - Project ID 390523135).

References

Artiga, M. (2016), ‘New perspectives on artifactual and biological functions’, *Applied Ontology* **11**, 89–102.

Boorse, C. (1976), ‘Wright on functions’, *The Philosophical Review* **85**(1), 70–86.

Carandini, M. & Heeger, D. J. (2012), ‘Normalization as a canonical neural computation’, *Nature Neuroscience* **13**, 51–62.

Chakravartty, A. (2010), ‘Perspectivism, inconsistent models, and contrastive explanation’, *Studies in History and Philosophy of Science* **41**, 405–412.

Chalmers, D. J. (2011), ‘A computational foundation for the study of cognition’, *Journal of Cognitive Science* **12**(4), 323–357.

Chalmers, D. J. (2012), ‘The varieties of computation: a reply’, *Journal of Cognitive Science* **13**, 211–248.

Coelho Mollo, D. (2018), ‘Functional individuation, mechanistic implementation: the proper way of seeing the mechanistic view of concrete computation’, *Synthese* **195**, 3477–3497.

Coelho Mollo, D. (2019), ‘Are there teleological functions to compute?’, *Philosophy of Science* **86**(3), 431–452.

Colombo, M. (forthcoming), (Mis)computation in computational psychiatry, *in* F. Calzavarini & M. Viola, eds, ‘New Challenges in Philosophy of Neuroscience’, Springer.

Colombo, M., Hartmann, S. & van Iersel, R. (2015), ‘Models, mechanisms, and coherence’, *The British Journal for the Philosophy of Science* **66**(1), 181–212.

Craver, C. F. (2013), Functions and mechanisms: a perspectivalist view, *in* P. Huneman, ed., ‘Functions: Selection and Mechanisms’, Springer.

Dewhurst, J. (2018*a*), ‘Computing mechanisms without proper functions’, *Minds & Machines* **28**(3), 569–588.

Dewhurst, J. (2018*b*), ‘Individuation without representation’, *British Journal for the Philosophy of Science* **69**(1), 103–116.

Fodor, J. A. (1985), ‘Fodor’s guide to mental representation: The intelligent auntie’s vade-mecum’, *Mind* **XCIV**(373), 76–100.

Fresco, N. (2014), *Physical Computation and Cognitive Science*, Springer.

Fresco, N. & Milkowski, M. (forthcoming), ‘Mechanistic computational individuation without biting the bullet’, *British Journal for the Philosophy of Science* .

Fresco, N. & Primiero, G. (2013), ‘Miscomputation’, *Philosophy and Tech-

nology **26**(3), 253–272.

Garson, J. (2011), 'Selected effects and causal role functions in the brain: the case for an etiological approach to neuroscience', *Biology and Philosophy* **26**(4), 547–565.

Garson, J. (2013), 'The functional sense of mechanism', *Philosophy of Science* **80**, 317–333.

Giere, R. N. (2006), *Scientific Perspectivism*, Chicago University Press.

Godfrey-Smith, P. (1993), 'Functions: Consensus without unity', *Pacific Philosophical Quarterly* **74**(3), 196–208.

Griffiths, P. E. (1993), 'Functional analysis and proper functions', *British Journal for the Philosophy of Science* **44**, 409–22.

Haimovici, S. (2013), 'A problem for the mechanistic account of computation', *Journal of Cognitive Science* **14**, 151–181.

Haugeland, J. (1985), *Artificial Intelligence: the very idea*, MIT Press.

Huneman, P., ed. (2013), *Functions: selection and mechanisms*, Springer.

Ismael, J. (2016), 'How do causes depend on us? The many faces of perspectivalism', *Synthese* **193**, 245–267.

Khuu, S. K., Honson, V. J. & Challinor, K. L. (2016), 'The influence of cast shadows on the detection of three-dimensional curved contour structure', *Perception* **45**(4), 425–442.

Kriegeskorte, N. & Douglas, P. K. (2018), 'Cognitive computational neuroscience', *Nature Neuroscience* **21**, 1148–1160.

Lehky, S. R. & Sejnowski, T. J. (1988), 'Network model of shape-from-shading: neural function arises from both receptive and projective fields', *Nature* **333**, 452–454.

Machamer, P., Darden, L. & Craver, C. (2000), 'Thinking about mechanisms', *Philosophy of Science* **67**, 1–25.

Maley, C. J. & Piccinini, G. (2017), A unified mechanistic account of teleological functions for psychology and neuroscience, *in* D. M. Kaplan, ed., 'Explanation and Integration in Mind and Brain Science', Oxford University Press, pp. 236–254.

Massimi, M. (2018), 'Four kinds of perspectival truth', *Philosophy and Phenomenological Research* **XCVI**(2), 342–359.

Milkowski, M. (2013), *Explaining the Computational Mind*, The MIT Press.

Millikan, R. G. (1984), *Language, Thought, and Other Biological Categories: New Foundations for Realism*, The MIT Press.

Morrison, M. (2011), 'One phenomenon, many models: Inconsistency and complementarity', *Studies in History and Philosophy of Science* **42**, 342–351.

Neander, K. (1991), 'Functions as selected effects: The conceptual analyst's defence', *Philosophy of Science* **58**(2), 168–184.

Neander, K. (2017), *A Mark of the Mental: in defense of informational semantics*, The MIT Press.

Piccinini, G. (2015), *Physical Computation: a Mechanistic Account*, Oxford University Press.

Putnam, H. (1988), *Representation and Reality*, MIT Press.

Schweizer, P. (2014), Algorithms implemented in space and time, *in* 'Selected Papers from the 50th Anniversary Convention of the AISB', pp. 128–136.

Schweizer, P. (2016), In what sense does the brain compute?, *in* V. C. Müller, ed., 'Computing and Philosophy', Vol. 375 of *Synthese Library*, Springer, pp. 63–79.

Schweizer, P. (2019), Computation in physical systems: a normative mapping account, *in* D. Berkich & M. V. d'Alfonso, eds, 'On the Cognitive, Ethical, and Scientific Dimensions of Artificial Intelligence - Themes from IACAP 2016', Springer, pp. 27–47.

Searle, J. R. (1992), *The Rediscovery of the Mind*, MIT Press.

Shagrir, O. (1999), 'What is computer science about?', *The Monist* **82**(1), 131–149.

Sprevak, M. (2013), 'Fictionalism about neural representations', *The Monist* **96**, 539–560.

Sprevak, M. (2018), Triviality arguments about computational implementation, *in* M. Sprevak & M. Colombo, eds, 'Routledge Handbook of the Computational Mind', Routledge, pp. 175–91.

Tucker, C. (forthcoming), 'How to explain miscomputation', *Philosophers' Imprint* .

Van Gelder, T. (1995), 'What might cognition be, if not computation?', *The Journal of Philosophy* **92**(7), 345–381.

Wang, X.-J. & Krystal, J. H. (2014), 'Computational psychiatry', *Neuron* **84**, 638–654.