

THE COGNITIVE BASIS OF COMPUTATION

Putting computation in its place

Daniel D. Hutto, Erik Myin, Anco Peeters and Farid Zahnoun

Computing is normally done by writing certain symbols on paper.

(Alan Turing, 1936, p. 249)

Use of computational models and talk of computation is rife in explanations of cognition. In philosophical hands, this anodyne observation about explanation is transformed when it is augmented by the claims that computational processes are metaphysically real processes that are either necessary or sufficient for cognition. Whether advanced in its weaker, necessity, or stronger, sufficiency guise, the underlying idea is that computation is *a*, if not *the*, explanatory basis for cognition. Call this underlying idea the Computational Basis of Cognition thesis, or CBC.¹

The CBC maintains that if biological minds or artificial systems are capable of cognition then they must, in one way or another, compute. The CBC motto is: no cognition without computation. With respect to the special case of putative brain-based, neural computations, advocates of CBC hold that such computations form the basis of, and explain, a wide range of cognitive operations. These not only include perception, attention, language-processing, and reasoning but also canonical acts of overt computation – like adding fractions, solving equations, generating proofs – that feature in specialized symbol-manipulating, rule-based practices. Notably in the latter cases, the direction of explanation runs from the covert to the overt in that overt computational processes are taken to be explained by covert computational processes and not the other way around.

The CBC can be supported in various ways, employing different theories of computation. In Section 1, we introduce representationalist theories of computation that might be used to ground the CBC, identify what motivates them, and raise concerns about such theories. Sections 2 and 3 focus on the prospect of grounding the CBC in non-representational theories of computation. Ultimately, we question the ability of such theories to deliver a metaphysically robust, naturalistic account of computation of the sort needed to support the CBC.² In the final section, we articulate the alternative possibility to the CBC – namely, that computation may depend on semantically-laden cognition and not the other way around. We put forward a reversed rival of the CBC – one that, we argue, avoids the problems encountered by both representational and non-representational causal-mechanistic theories of computation and which is consistent with the known facts about how minds and brains work.

1 The computational basis of cognition: representationalist theories

Some defenders of the CBC embrace representational theories of computation. Representational theories of computation propose that computations are always and everywhere operations over symbols and that symbols have both representational and syntactic properties essentially (Salisbury and Schneider, this volume). In promoting this style of theory, Fodor, perhaps the staunchest of its advocates, maintains that “all symbol tokens have intentional contents and ... syntactic structure – constituent structure in particular – that’s appropriate to the content they have” (Fodor, 1987, pp. 135–137; see also Fodor 1990, p. 167; 1975, p. 198).

There are two interlocking assumptions at the heart of representational theories of computation. The first assumption is that symbols are taken to be individuated partly by their syntactic structure. It is their syntactic structure that links them inferentially to other symbols. Symbols are assumed to have a kind of ‘shape’, analogous to the shapes of the syntactical forms of natural language sentence tokens, where such shapes have structural properties that cannot be understood solely in terms of their physical properties. Nevertheless, it is the fact that the syntactic properties of symbols can be implemented concretely in physical systems that makes computational processes mechanically possible.

The second assumption is that symbols are also taken to be partly individuated by what they are about – namely, what they denote or refer to. Crucially, according to the ‘received view’, as Sprevak (2010) dubs it, a symbol’s representational properties determine which computations, if any, are taking place. For theorists attracted to this view, the fact that symbols have the representational properties that they do solves the otherwise intractable problem of computational individuation; namely, determining whether a given process is a computational process by specifying which function or rule is being carried out.

To borrow an example from Sprevak (2010), consider an electrical system which receives either 0V or 5V at its two input nodes. The system will output 5V only if its two input nodes receive 5V, as specified in Table 20.1. At first sight, the system may be thought to implement a classic AND-gate. However, this assumes that the 5V output has a ‘1’ or ‘true’ value. But 5V could have a ‘0’ or ‘false’ value, in which case the system would be implementing an OR-gate.

Apparently, the non-semantic properties of the input–output patterns alone are insufficient to determine which of the two values are in play and hence which computational function is being implemented. Yet the problem is overcome if, as Sprevak (2010) claims, “Appeal to representation allows us to decide between these two options ... [and that] the difference between an implementation of an AND gate and an OR gate is a difference in representational content” (Sprevak 2010, p. 269).

This observation motivates the idea that representational contents are needed to specify which computational functions are being carried out. Generalizing, the lesson to be drawn from this case is that representational contents are needed to do the heavy lifting in fixing

Table 20.1 A specification of an electrical system

<i>IN1</i>	<i>IN2</i>	<i>OUT</i>
0V	0V	0V
0V	5V	0V
5V	0V	0V
5V	5V	5V

computational identities and that “the functions a system computes are always characterized in semantic terms” (Shagrir, 2001, p. 382).

In sum, the problem of computational individuation can apparently be dealt with if it is assumed that “representational content plays an essential role in the individuation of states and processes into computational types” (Shagrir, 2006, p. 393). Realistically construed, as O’Brien and Opie (2009) put it, computation is, at least partly, dependent upon and “governed by the contents of the representations it implicates” (p. 53). The cost of accepting that computational states and input–output–functions are essentially, if only partially, individuated semantically is to accept that: “there is no computation without representation” (Fodor, 1981, p. 180).

Ultimately, the price of the representationalist solution to the problem of computational individuation may be too high for defenders of the CBC who also advocate explanatory naturalism. For to make a representational theory of computation work we need a theory that tells us “what counts as representation” and “what gives representations their content” (Piccinini, 2015, p. 29; see also Chalmers, 2011; 2012).³

Naturalists who subscribe to a representational theory of computation and hope to use it to mount a credible defense of the CBC are, in the end, obliged to supply a grounding theory of content. In line with the CBC such a theory would need to explain how computational vehicles gain their contents without making appeal to the norms and rules supplied by socio-cultural practices. The reasoning is straightforward: if, as the CBC assumes, computation is required for cognition, and if cognition is required for sociocultural practices, then, by implication, computation is required for sociocultural practices. This places an important constraint on the CBC: if there can be no computation without representation then the representations in question must be accounted for independently of and prior to the emergence of social-cultural practices.

Yet despite many dedicated efforts, we currently lack a tenable naturalized theory of content – given in causal, informational, or biological terms or some combination thereof – that satisfies the demands of the CBC. Without such a theory as a principled means of allocating contents to vehicles, representational theories of computation remain, at best, programmatic and promissory. Certainly, for anyone attracted to explanatory naturalism, such theories do not supply a secure foundation for the CBC here and now.⁴

In what other ways might the CBC be defended, assuming that these concerns constitute reasons to steer clear of representationalist theories of computation?

Several philosophers have proposed that a tenable non-representational notion of computation is already well within our reach. There are various theories on the market that attempt to define computations solely in formal, structural, or mechanical terms (e.g. Chalmers, 2011; Miłkowski, 2013; Piccinini, 2015). As their collective name indicates, all of the theories in this family seek to demonstrate how “computation can be fully individuated without appeal to semantic content” and how “there is computation without representation” (Piccinini, 2015, p. 33).⁵ As such, should any of these theories prove workable it would supply an account of computation that avoids the need to naturalize semantic content.

Ultimately, with respect to the CBC, more is required. As far as securing the CBC goes, the crucial test of any tenable non-semantic theory of computation is whether it articulates a notion of computation that will prove foundational in the sciences of the mind. Thus, the pivotal question is: are there any tenable non-semantic theories of computation and, if so, can they play such a role?

In the end, there are reasons to doubt that non-semantic theories can provide an account of computation needed to do the sort of explanatory work required for securing the CBC. A full survey of such theories and their potential to secure the CBC is beyond the scope of this

chapter. Still, it is possible to highlight the main sort of challenges this class of theories faces by focusing on two representative samples.

2 The computational basis of cognition: a non-semantic, functional theory

Drawing on a long tradition inspired by the properties of Turing machines, Chalmers (2011; 2012) offers a functionalist account of computation – one that is meant to capture the core understanding of computation as it figures in the formal theories of computer scientists. The central assumption of this functionalist theory is that a computation is a formalism that specifies a system’s causal topology – namely, its fine-grained organizational structure – by specifying the system’s inputs, outputs, internal states and their transitions.⁶

To this account of what computations are, Chalmers adds a general theory of what it takes for a physical system to implement a computation. On a rough-and-ready rendering, such implementation occurs when the causal structure of some concrete physical system formally mirrors the structure of the computation. It follows from this theory of computational implementation that computations abound in nature since any appropriately organized physical system will implement at least one computation.

With respect to the CBC, Chalmers (2011; 2012) makes an important claim about cognition – one that goes beyond his general theories of computation and computational implementation. He holds that, perhaps with some exceptions, cognitive properties and processes are maximally indifferent to the material substrates of the systems in which they are implemented: they are distinguished in being organizationally invariant.⁷ A property or process is organizationally invariant if it can be implemented concretely in some physical system merely by implementing its causal topology.

Chalmers denies that every physical process has this feature. Digestion, he contends, does not. This is because he assumes that digestion requires particular physio-chemical properties. Changing the material substrates in certain respects, but retaining the causal topology, does not guarantee digestion. For cognitive processes on the other hand, merely retaining its causal topology and implementing it by any material means secures that the implementation is a *bona fide* instance of cognition.

There are two problems in attempting to secure the CBC by appeal to the supposed organizational invariance of cognition. First, and this is the principal concern, even if it turns out that cognition has the property of being organizationally invariant, the fact that cognition is computational in Chalmers’ general sense would not explain why cognition has this special feature. Chalmers simply highlights that, on his account, cognition is special because it can be implemented in any system that preserves the relevant causal topology. Yet his appeal to cognition’s alleged computational nature would not explain this special status, even if it turns out to be analytically true.

Second, it is contentious whether cognition actually differs in kind from other processes, like digestion, in being maximally indifferent to its material substrates. As far as anyone knows cognition may depend to a much greater extent on its material substrates than Chalmers imagines. After all, it is no accident that sensory systems have the particular material properties that they have. This is because having such properties appears to be required if they are to fulfil their cognitive functions. As Dennett (1997) observes, “in order to detect light ... you need something photosensitive” (p. 97). This is the case not just for natural but also for artificial eyes. Even though there is still some degree of flexibility in how visual perception might be achieved, this fact places significant limits on what materials might be used for visual systems to get their cognitive work done.

The story is not importantly different when it comes to the cognitive contributions of brains:

The recent history of neuroscience can be seen as a series of triumphs for the lovers of detail ... the specific geometry of the connectivity matters ... the location of specific neurotransmitters and their effects matters ... the architecture matters ... the fine temporal rhythms of spiking patterns matter, and so on.

(Dennett, 2005, p. 19)⁸

It may be that for some cognitive operations only certain physical properties matter, such as the timing of neuronal spiking patterns. Still, it would not follow that such properties are abstract and substrate-indifferent as opposed to being concrete and substrate-sensitive. As Polger and Shapiro (2016) emphasize in resisting the former interpretation, “the frequency of the spike train of a neuron or neural assembly ... is a property of neurons as neurons, not just as implementers of some supraneural process” (p. 164).

In short, it would be hard to deny that cognitive processes depend on particular materials despite exhibiting varying degrees of substrate-neutrality.⁹ What is not established is that cognitive processes are maximally substrate-neutral such that it is possible to re-create all their relevant causal patterns in alternative media. Thus, on the question of whether cognition exhibits organizational invariance the jury is still out.

What if the claim that cognition is organizationally invariant fails to garner empirical support? Is there still enough strength in an unvarnished Chalmers-style functionalist theory of computation to establish the CBC? In one sense, it might appear so. After all, the cornerstone assumption of Chalmers’ theory is that anything with a causal structure that can be specified by means of a computational formalism implements a computation.

Towl (2011) complains that this feature of a Chalmers-style theory of computation makes it overly permissive: namely, it threatens to trivialize the notion of computation. He asks us to consider a game at a pool table. It is easy enough to characterize the physical events of the game – such as the movement of the balls – as implementing a form of vector addition in terms of direction and velocity. But what explanatory advantage is conferred by treating the activity of the balls as implementing computations? What is explained by supposing the balls and the pool table are in fact computing vector sums? As Towl stresses, the situation would be entirely different if we used the movements of the balls to compute sums or if we connected the movement of the balls to a specialized device that was dedicated to the purpose of computing sums.

This type of complaint does not trouble Chalmers’ general theory. He admits that his account of computation may fail to capture useful distinctions required for certain explanatory purposes. That is fine so long as it serves other needs. He is attracted to a pluralism that allows him to isolate the value of his general theory of computation to picking out and demarcating the formal subject matter that is of special interest to the computational sciences (Chalmers, 2012).¹⁰

Yet even if the pluralist reply holds up, Chalmers’ general theory of computation would at most answer a classificatory need: it would not have the explanatory punch needed to defend the CBC. After all, even if we suppose that brains, just like pool tables, implement computations in the way that Chalmers’ general theory assumes, we would need a story about how and why any such computations make a difference to and explain cognition.

In the end, if our understanding of the computational theory of cognition is based solely in a Chalmers-style general theory of computation then it loses “much of its explanatory force” (Piccinini, 2015, p. 55). Thus, the price of guaranteeing the CBC by appeal to Chalmers’ theory of computation is that the CBC is rendered trivially true but explanatorily hollow with respect to the needs of the sciences of the mind (O’Brien, 2011; Ritchie, 2011; Rescorla, 2012).

3 The computational basis of cognition: a non-semantic, mechanistic theory

Piccinini (2015) advances a mechanistic account of computation which is designed to overcome the sorts of problems faced by more liberal functionalist theories of computation. It operates with a generic definition of computation that restricts the class of physical computing systems to a sub-class of functional mechanisms. The central plank of this theory is that to qualify as a computing system a mechanism must have the function to manipulate medium-independent vehicles according to rules as one of its teleo-functions.

The key assumptions of this theory are as follows: A vehicle is understood as a variable – a state that can take different values and change over time – or a specific instance of such a variable (Piccinini, 2015, p. 121). Vehicles are manipulated “according to rules that are sensitive solely to differences between different portions (i.e., spatiotemporal parts) of the vehicles” (ibid., p. 121; see also Piccinini and Bahar, 2013, p. 458). Rules are here understood broadly and in non-representational terms: rules are simply input to output maps. Finally, and crucially, all concrete computations and their vehicles are deemed medium-independent because they can be described and defined “independently of the physical media that implement them” (Piccinini, 2015, p. 122).¹¹

In operating with a much more restrictive theory of what counts as a physical computation than its purely functionalist rival, Piccinini’s mechanistic theory demarcates computing systems from other sorts of functional devices in a way that avoids pan-computationalism. Consequently, by its lights, digestive systems and pool tables lack the special features just mentioned needed to qualify as computing systems. Moreover, with respect to the CBC, this theory looks, *prima facie*, far better placed than its functionalist rival to deliver the required explanatory goods.

There is one apparent obstacle to defending the CBC by appeal to a mechanistic theory of computation of this sort. It is that there are clear dissimilarities between what happens in brains and what happens in artifactual computers. Indeed, looking solely at the character of neural activity it has been observed that brains are not executing computations of any familiar kind – brains are not performing digital or analog computations. Summarizing an analysis of a wide range of findings, Piccinini and Bahar (2013) openly acknowledge this fact, reporting that, “in a nutshell, current evidence indicates that typical neural signals, such as spike trains ... are neither continuous signals nor strings of digits” (p. 477).

Can we infer from these observations that brains are not performing any kind of computation? No. Piccinini and Bahar (2013) conclude that brains are performing computations of a special variety, maintaining that neural computation happens in its own special way – namely that “neural computation is *sui generis*” (p. 477; see also Piccinini, 2015, p. 223). Of course, this inference is not obligatory. If the above evidence were all we had to go on then we would be equally justified in concluding that brains do not compute.

Why then suppose, in light of such findings and the constraints of a mechanistic theory of computation, that the neural processes that contribute to explaining cognition are computational? Piccinini and Bahar (2013) supply an argument based on the following assumptions: cognition involves information processing of a kind that requires the manipulation of “vehicles based on the information they carry rather than their specific physical properties” (p. 463). Therefore, cognition requires the manipulation of medium independent vehicles. Hence, the neural processes that contribute to cognition must involve the manipulation of medium-independent vehicles.¹²

Voltage changes in dendrites, neuronal spikes, neurotransmitters, and hormones are offered as prime examples of neurocomputational vehicles. Piccinini and Bahar (2013) hold that such neural events and entities qualify as medium-independent vehicles because the properties

which are relevant for their cognitive work – such as firing rates and patterns – can be defined in abstract terms. Thus, these authors claim, this makes such vehicles unlike the other, putatively more concrete properties of the neural systems that implement them.

It is questionable, however, that the neural events and processes that underpin cognition actually have the feature of being medium-independent. There is reason to doubt that neural events could contribute to cognitive work if that work really requires the concrete manipulation of medium-independent vehicles. The trouble is that if medium-independent vehicles are defined by their abstract properties then it is unclear how such vehicles could be concretely manipulated. Understanding how neural processes can be sensitive to concrete, medium-dependent properties presents no conceptual difficulty. By contrast, we have no conception of how concrete neural processes could causally manipulate abstract, medium-independent vehicles. Certainly, the defenders of the mechanistic theory of computation offer no account of how such manipulations might be achieved.

Again, there is no barrier to understanding how neural events can be sensitive to only specific aspects of a concrete structure. Nor is there a barrier to understanding how an analogue of that neural process could be sensitive to the same aspect of an analog structure in a materially different system. But that does not make either the imagined neural process or its analogue sensitive to a medium-independent property. Neither process is sensitive to what the other process is sensitive to. Rather, they are both sensitive to some aspect of physical structures that can be given a medium-independent description.

Polger and Shapiro (2016) diagnose the source of confusion that gives rise to belief in abstract medium independent vehicles, as one of conflating the abstractness which is a feature of computational models with features of “the processes being modelled” (p. 166). Elaborating, they observe that “the apparent medium independence of computational explanations owes to the fact that they model or describe their phenomena in topic-neutral or abstract ways rather than to the abstractness or multiple realizability of their objects” (ibid., p. 155).

In the end, it turns out that medium independence is not a property of physical token processes, but rather is a relational or comparative property of several processes. As a result, one can have medium-independent descriptions of processes – descriptions which abstract from certain substrate-related properties and mention properties which can be found in different substrates – but one cannot have concrete vehicles that are medium-independent.

Happily, the dimensions of variation in physical systems to which neural events are sensitive need not be construed as medium-independent vehicles. They may simply be dimensions of variation in the concrete properties of certain structures. Nor need the lawful changes involved in being sensitive to certain properties of such structures be thought of in terms of the rule-bound manipulation of medium-independent vehicles. They might simply be systematic changes that conform to specific patterns.

In sum, these considerations cast serious doubt on the possibility of employing a non-semantic mechanistic theory of computation to support the CBC.

4 The cognitive basis of computation: a sociocultural theory

As the preceding analysis reveals, there are serious problems with the most promising existing proposals for securing the CBC. As things stand, there is no compelling evidence or theoretical argument for supposing that computation is the, or even an, explanatory basis of cognition.

Where in the world, then, do we find computations and how do they relate to minds? There is another possibility to consider – one left hanging at the end of the first section: namely, that computation may depend on cognitive activity and not the other way around. The kind of

cognition in question, we propose, is that which only arises within and is integrally bound up with specific sociocultural normative practices.

In locating computations in nature, we seek to revive the original model of a computational system (see, Copeland and Proudfoot, this volume; Isaac, this volume) which was that of “a person – a mathematician or logician manipulating symbols with hands and eyes, and pen and paper (The word ‘computer’ originally meant ‘one who computes’)” (Thompson, 2007, p. 7). Accordingly, in the originary case computing first arises along with the emergence of “a sophisticated form of human activity” (ibid., p. 7).

Computation originally consists of symbol manipulating operations carried out by people. Sociocultural practices make it the case that certain operations with symbols are properly identifiable and individuated as computations: the reason is because it is only within such practices that computational operations and manipulations have a home. Such computations are semantically laden, in the sense that statements which express particular computational operations, such as the result of calculating a derivative function, are true or false. The surrounding context and practices of such manipulations determines whether a given manipulation of symbols is an instance of computing or not. This is because both the current use and the larger history of a person or system determines whether the manipulation forms part of, say, a particular computational operation, some other computational operation, or none at all. Borrowing an example from Michael Rescorla, it is the surrounding history and practices that determine whether a child in a contemporary context, while performing an arithmetic operation over numerals, is computing in the decimal system and not in some other system like base-13 (Rescorla, 2013).

Sociocultural practices for structurally manipulating tokens in specific ways that accord with an established practice is plausibly not only the basis for how human beings compute, it is likewise the basis for artificial forms of computation. We often rely on artifacts and artificial systems to compute with and for us. We can compute by writing with chalk on blackboard; by moving the beads of an abacus; or by pressing the keys of a calculator. In such circumstances, we compute with chalk and board, with the abacus, or with the calculator. Yet, focusing on the last case, we not only compute with calculators, we also say that calculators compute. The only relevant difference is that computing with an abacus requires moving the beads around, while computing with a calculator requires pressing some keys and then letting the mechanics of the machine take care of the rest.

Importantly, when we construct artificial computing devices we do so by relying on, and re-arranging concrete physical materials, so that they acquire a structure that suits our goals. There is no reason to assume that, before these materials or processes are put to computational use by us, they already compute. In other words, constructing computers consists in transforming material devices and processes that do not compute into devices and processes that we can compute with. The computational properties of these devices depend on the surrounding sociocultural activities of which they become part. Accordingly, we can think of Turing machines that manipulate meaningless strings as simply not computing until those strings are put to use for specific purposes.

Perhaps some of the processes occurring in brains are, at some level of abstraction, similar to the kinds of processes found in our computational artifacts. But that does not imply that the brain computes, only that we can draw analogies between these two kinds of processes. Neither does the fact that people can compute ‘in their heads’, without engaging in overt manipulation of symbols, show that in such cases the brain computes. Even in these cases, it is the person that computes. The fact that computing relies on, and would not be possible without the occurrence of specific brain processes does not entail that those brain processes themselves are computations.

If brains are not computing when contributing to cognition then what are they doing? Neural activity is sensitive to relations of covariance, and such sensitivity drives cascades of neural activity that influence and constrain organismic responsiveness. But such coordinated activity need not be thought of as processing information or, thereby, as a kind of computation. That assumption is not necessary to explain the work that brains do in enabling organisms to “get a grip on the patterns that matter for the interactions that matter” (Clark, 2016, p. 294). Well-calibrated neural activity can systematically influence and constrain organismic responding, and even maintain connections with specific worldly features without the brain engaging in any computations. In other words, neurodynamics can be, and apparently should be, conceived of in terms of coordinated cascades rather than in terms of information processing computations (Hutto and Myin, 2017, epilogue).

Why take this sociocultural proposal about the basis of computation seriously? As we have seen there are inherent difficulties in supposing that computation arises in nature independently of and prior to socioculturally based practices of someone or something computing for a purpose.

How, on our account, does our sociocultural account of computation relate to cognition? If the analysis and arguments of this chapter hold up then we have reason to try to invert the explanatory order proposed by the Computational Basis of Cognition thesis. We must reverse the polarity of standard thinking on this topic, and ask how it is possible that computation, natural and artificial, might be based in cognition and not the other way around.

If specific sociocultural practices are a necessary and sufficient explanatory basis of computation, and those practices are themselves cognitively based, then it follows that computation is also cognitively based. Of course, the cognitive basis need not itself be representational (see Hutto and Myin, 2013; 2017; Hutto and Satne, 2015); and if we are correct, on pain of circularity, it cannot be computational either.

Acknowledgments

Daniel Hutto thanks the Australian Research Council, Discovery Project DP170102987 ‘Mind in Skilled Performance’, and Erik Myin and Farid Zahnoun thank the Research Foundation Flanders (FWO), projects G048714N ‘Offline Cognition’ and G0C7315N ‘Getting Real about Words and Numbers’, for funding that enabled the completion of the primary research informing this chapter. We are also grateful to Mark Sprevak and Matteo Colombo for their invitation to contribute, and for their excellent and helpful feedback.

Notes

- 1 Notably the CBC, in either its necessity or sufficiency variant, is a much stronger thesis than the thesis “that concept of computation lies at the very foundation of cognitive science” (O’Brien, 2011, p. 381). Thus a prominent neurocentric version of the CBC espouses that: “brains perform computations and neural computations explain cognition” (Piccinini, 2015, p. 207). Piccinini and Bahar (2013) trace an industrial strength, neural variant of the CBC – that neural activity simply *is* computation – back to McCulloch and Pitts (1943). Whether articulated in stronger or weaker form, the CBC has more or less enjoyed the status of the received view in the sciences of the mind ever since the advent of the cognitive revolution: see Piccinini (2015, p. 207) for a long list of those who have defended this idea in some shape or form since the 1970s forward. Indeed, support for the neural variant of the CBC runs so deep that, as Piccinini and Bahar (2013) report, many cognitive scientists even “consider it commonsensical to say that neural activity is computation and that computation explains cognition” (p. 454).
- 2 For reasons of space we do not discuss other, less widely endorsed theories of computation. See Piccinini (2015, chs. 2–4) for a more systematic review of other positions and the problems they face.

- 3 Representational theories of computation are accused of having feet of clay. Those at the vanguard of these debates have observed that, “the notion of semantic content is so ill-understood that it desperately needs a foundation itself” (Chalmers, 2011, p. 334).
- 4 As long as semantic or representational content is understood in terms of having satisfaction conditions of some kind – for example, truth or accuracy conditions – then there are reasons to think that no naturalistic theory of content is anywhere in sight. To supply such a theory would require overcoming the Hard Problem of Content (Hutto and Myin, 2013; 2017). Until that problem is dealt with, there is no gain in appealing to semantic or contentful properties that allegedly permeate and individuate computational processes.
- 5 Importantly, non-semantic accounts of computation can allow that computations can involve the manipulation of vehicles bearing representational contents. This can be the case, according to such theories, just so long as representational contents are not taken to be essential to the existence of computational processes (Chalmers, 2011; Miłkowski, 2013; Piccinini, 2015).
- 6 According to Chalmers (2011) a causal topology is “the abstract causal organization of the system: that is, the pattern of interaction among parts of the system, abstracted away from the make-up of individual parts and from the way the causal connections are implemented” (p. 337).
- 7 Chalmers (2011) observes that not all aspects of cognition will be organizationally invariant: any aspect of cognition that partly depends on the actual make-up of the environment will not. He gives knowledge and belief as examples, on the assumption that their contents are fixed by external factors. Famously, if content externalism holds, whether one has a belief about water or not depends on the actual physio-chemical make-up of the relevant substances in the world one occupies.
- 8 Importantly, on this score Chalmers (2012) acknowledges that “locations, velocities, relative distances and angles are certainly not organizational invariants: systems with same causal topology might involve quite different locations, velocities and so on” (p. 216).
- 9 Even Chalmers allows that digestion can survive some changes to its physio-chemical substrate so long as the relevant causal patterns are preserved (2011, p. 338). Thus digestion may be at one end of the substrate-neutrality spectrum and certain cognitive processes at the other.
- 10 It is far from obvious that Chalmers is right on this score – viz., that his general theory adequately captures such scientific commitments. There is a great deal of disagreement in the field about which notion of computation is in fact deployed in computability theory and computer science (Piccinini, 2008, p. 6; Rescorla, 2017, p. 8).
- 11 Piccinini cites Garson as the inspiration for his strong construal of medium independence (see Garson, 2003). An earlier formulation of medium independence can be found in Haugeland (1989) when he speaks of formal systems being realized in “any number of different media” (p. 58).
- 12 See Hutto and Myin (2013; 2017) for reflections about the nature of information that support the idea that cognition involves medium-dependent information sensitivity as opposed to medium-independent information processing.

References

- Chalmers, D. (2011) ‘A Computational Foundation for the Study of Cognition’, *Journal of Cognitive Science*, 12 (4), pp. 323–357.
- Chalmers, D. (2012) ‘The Varieties of Computation: A Reply’, *Journal of Cognitive Science*, 13 (3), pp. 211–224.
- Clark, A. (2016) *Surfing Uncertainty: Prediction, Action and the Embodied Mind*. Oxford: Oxford University Press.
- Dennett, D.C. (1997) *Kinds of Minds: Towards an Understanding of Consciousness*. London: Phoenix.
- Dennett, D.C. (2005). *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness*. Cambridge, MA: MIT Press.
- Fodor, J.A. (1975) *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fodor, J.A. (1981) *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Fodor, J.A. (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, J.A. (1990) *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Garson, J. (2003) ‘The Introduction of Information into Neurobiology’, *Philosophy of Science*, 70 (5), pp. 926–936.
- Haugeland, J. (1989) *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.

- Hutto, D.D. and Myin, E. (2013) *Radicalizing Enactivism: Basic Minds without Content*. Cambridge, MA: MIT Press.
- Hutto, D.D. and Myin, E. (2017) *Evolving Enactivism: Basic Minds Meet Content*. Cambridge, MA: MIT Press.
- Hutto, D.D. and Satne, G. (2015) 'The Natural Origins of Content', *Philosophia*, 43 (3), pp. 521–536.
- McCulloch, W.S. and Pitts, W. (1943) 'A Logical Calculus of the Ideas Immanent in Nervous Activity', *The Bulletin of Mathematical Biophysics*, 5 (4), 115–133.
- Miłkowski, M. (2013) *Explaining the Computational Mind*. Cambridge, MA: MIT Press.
- O'Brien, G. (2011) 'Defending the Semantic Conception of Computation in Cognitive Science', *Journal of Cognitive Science*, 12 (4), pp. 381–399.
- O'Brien, G and Opie, J. (2009) 'The Role of Representation in Computation', *Cognitive Processing*, 10 (1), pp. 53–62.
- Piccinini, G. (2008) 'Computation without Representation', *Philosophical Studies*, 137 (2), pp. 205–241.
- Piccinini, G. (2015) *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- Piccinini, G. and Bahar, S. (2013) 'Neural Computation and the Computational Theory of Cognition', *Cognitive Science*, 37 (3), pp. 453–488.
- Polger, T.W. and Shapiro, L.A. (2016) *The Multiple Realization Book*. Oxford: Oxford University Press.
- Rescorla, M. (2012) 'How to Integrate Representation into Computational Modelling, and Why We Should', *Journal of Cognitive Science*, 13 (1), pp. 1–38.
- Rescorla, M. (2013) 'Against Structuralist Theories of Computational Implementation', *The British Journal for the Philosophy of Science*, 64 (4), pp. 681–707.
- Rescorla, M. (2017) 'Levels of Computational Explanation', in Powers, T.M. (ed.) *Philosophy and Computing*, vol. 128. Cham: Springer, pp. 5–28.
- Ritchie, J.B. (2011) 'Chalmers on Implementation and Computational Sufficiency', *Journal of Cognitive Science*, 12 (4), pp. 401–417.
- Shagrir, O. (2001) 'Content, Computation and Externalism', *Mind*, 110 (438), pp. 369–400.
- Shagrir, O. (2006) 'Why We View the Brain as a Computer', *Synthese*, 153, pp. 393–416.
- Sprevak, M. (2010) 'Computation, Individuation, and the Received View on Representation', *Studies in History and Philosophy of Science*, 41, pp. 260–270.
- Thompson, E. (2007) *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
- Towl, B.N. (2011) 'Home, Pause, or Break: A Critique of Chalmers on Implementation', *Journal of Cognitive Science*, 12 (4), pp. 419–433.
- Turing, A.M. (1936) 'On Computable Numbers, with an Application to the Entscheidungsproblem', *Proceedings of the London Mathematical Society*, 2nd series, 42 (1), pp. 230–265.