Objections to Computationalism. A Short Survey

Marcin Miłkowski (mmilkows@ifispan.waw.pl)
Instytut Filozofii i Socjologii PAN, ul. Nowy Świat 72,
00-330 Warszawa, Poland

Abstract

In this paper, I review the objections against the claim that brains are computers, or, to be precise, information-processing mechanisms. By showing that practically all the popular objections are based on uncharitable (or simply incorrect) interpretations of the claim, I argue that the claim is likely to be true, relevant to contemporary cognitive (neuro)science, and non-trivial.

Keywords: computationalism; computational theory of mind; representation; computation; modeling

Computationalism and Objections

The computational theory of mind, or computationalism, has been fruitful in cognitive research. The main tenet of the computational theory of mind is that the brain is a kind of information-processing mechanism, and that information-processing is necessary for cognition; it is non-trivial and is generally accepted in cognitive science. The positive view will not be developed here, in particular the account of physical computation, because it has already been elucidated in book-length accounts (Fresco, 2014; Miłkowski, 2013; Piccinini, 2015). Instead, a review of objections is offered here, as no comprehensive survey is available.

The survey suggests that the majority of objections fail just because they make computationalism a straw man. Some of them, however, have shown that stronger versions of the computational theory of mind are untenable, as well. Historically, they have helped to shape the theory and methodology of computational modeling. In particular, a number of objections show that cognitive systems are not *only* computers, or that computation is not the sole condition of cognition; no objection, however, establishes that there might be cognition without computation.

Computer metaphor is just a metaphor

Computational descriptions are sometimes described as *a computer metaphor* (cf., e.g., Ekman, 2003; Karl, 2012, p. 2101). The use of the term suggests that the proposed description is rough and highly idealized, and cannot be treated *literally*. However, by using the term, others suggest that no computational model may be treated seriously; all are mere metaphors (Daugman, 1990).

A defender of computationalism might concede this and weaken their position. But the position is also tenable in the stronger version. This is because computer metaphors cannot really be tested and rejected, whereas computational models can. For this reason, in this paper, I will adopt, along with other theorists (Newell & Simon, 1972, p. 5; Pylyshyn,

1984, pp. xiv–xvi), a stronger version of computationalism, which claims that cognition literally involves computation.

Software is not in the head

This objection is that there is no simple way to understand the notions of *software* and *hardware* as applied to biological brains. But the software/hardware distinction, popular as in the slogan "the mind to the brain is like the software to hardware" (Block, 1995; Piccinini, 2010), need not be applicable to brains at all for computationalism to be true. There are non-program-controllable computers: they do not load programs from external memory to internal memory in order to execute them. A mundane example of such a computer is a logical AND gate. In other words, while it may be interesting to inquire whether there is software in the brain, even if there were none, computationalism could still be true.

Computers are just for number-crunching

Another intuitive objection, already stated (and defeated) in the 1950s, is that brains are not engaged in number-crunching, while computers compute over numbers. But if this is all computers do, then they don't control missiles or send documents to printers. After all, printing is not *just* number crunching. The objection rests therefore on a mistaken assumption that computers can only compute numerical functions. Computer functions can be defined not only of integer numbers but also of arbitrary symbols (Newell, 1980), and as physical mechanisms, computers can also control other physical processes.

Computers are abstract entities

Some claim that because symbols in computers are, in some sense, abstract and formal, computers—or at least computer programs—are abstract as well (Barrett, 2015; Barrett, Pollet, & Stulp, 2014; Lakoff, 1987). In other words, the opponents of computationalism claim that it implies ontological dualism (Searle, 1990). However, computers are physical mechanisms, and they can be broken, set on fire *etc*. These things may be difficult to accomplish with a collection of abstract entities. Computers are not *just* symbol-manipulators. They do things, and some of the things computers do are not computational. In this minimal sense, computers are physically embodied, not unlike mammal brains. It is, however, a completely different matter whether the symbols in computers mean anything.

People are organisms, computers are not

Barrett (2015), among others, also presses the point that people are organisms. It's trivially true but irrelevant:

physical computers are physical, and they may be built in various ways. A computer may be built of DNA strands (Zauner & Conrad, 1996), so why claim that it's metaphysically impossible to have a biological computer?

Symbols in computers mean nothing

One of the most powerful objections formulated against the possibility of Artificial Intelligence is associated with John Searle's Chinese Room thought experiment (Searle, 1980). Searle claimed to show that running a computer program is not sufficient for semantic properties to arise, and this was in clear contradiction to what was advanced by proponents of Artificial Intelligence, who assumed that it was sufficient to simulate the syntactic structure of representations for the semantic properties to appear. As John Haugeland quipped: "if you take care of syntax, the semantics will take care of itself" (Haugeland, 1985, p. 106). But Searle replied: one can easily imagine a person with a special set of instructions in English who could manipulate Chinese symbols and answer questions in Chinese without understanding it at all. Hence, understanding is not reducible to syntactic manipulation. While the discussion around this thought experiment is hardly conclusive (Preston & Bishop, 2002), the problem was soon reformulated by Stevan Harnad (1990) as "the symbol grounding problem" (SGP): How can symbols in computational machines mean anything?

If the SGP makes sense, then one cannot simply assume that symbols in computers mean something just by being parts of computers, or at least they cannot mean anything *outside* the computer so easily (even if they contain instructional information (Fresco & Wolf, 2013)). Representational properties do not necessarily exist in physical computational mechanisms (Egan, 1995; Fresco, 2010; Miłkowski, 2013; Piccinini, 2008). So, even if Searle is right and there is no semantics in computers, the brain might still be a computer, as computers need no semantics to be computers. Perhaps something additional to computation is required for semantics.

There *is* an important connection between the computational theory of mind and the representational account of cognition: they are more attractive when both are embraced. Cognitive science frequently explains cognitive phenomena by referring to semantic properties of mechanisms capable of information-processing (Shagrir, 2010a). Brains are assumed to model reality, and these models can be utilized in computations. While this seems plausible to many, one can remain computationalist without assuming representationalism (the claim that cognition requires cognitive representation). At the same time, a plausible account of cognitive representation cannot be couched merely in computational terms as long as one assumes that the symbol grounding problem makes sense at least for some computers. To make the account plausible, most theorists appeal to notions of teleological function and semantic information (Bickhard, 2008; Cummins & Roth, 2012; Dretske, 1986; Millikan, 1984), which are not

technical terms of computability theory, neither can they be reduced to such. However, processing of semantic information is still processing of information; hence, computation is necessary for manipulation of cognitive representation.

Computationalism was strongly connected to cognitive representations by the fact that it offered a solution to the problem of what makes meaning causally relevant. Many theorists claim that because the syntax in computer programs is causally relevant (or efficacious), so is the meaning. While the wholesale reduction of meaning to syntax is implausible, the computational theory of mind makes it clear that the answer to the question includes the causal role of the syntax of computational vehicles. Still, the fact that it does not offer a naturalistic account of meaning is not an objection to computationalism itself. That would indeed be too much. At the same time, at least some naturalistic accounts, such as Millikan's and Dretske's, can be used to solve the SGP (see Miłkowski 2013, chap. 4).

Computers can only represent with all detail

The debate over meaning in computers and animals abounds in red herrings, however. One recent example is Robert Epstein's (2016) popular essay. His most striking mistake is the assumption that computers always represent everything with arbitrary accuracy. Epstein cites the example of how people remember a dollar bill, and assumes that computers would represent it in a photographic manner with all available detail. This is an obvious mistake: representation is useful mostly when it does not convey information about all properties of the represented target. If Epstein is correct, then there are no JPEG files in computers, as they are not accurate, because they are based on lossy compression. Moreover, no assumption of the computational theory of mind says that memory should be understood in terms of the von Neumann architecture, and it is controversial to suggest that it should (Gallistel & King, 2010).

People don't process information

Ecological psychologists stress that people do not process information, they just pick it up from the environment (cf. Chemero, 2003; Gibson, 1986). Thus, to understand this, one should make more explicit the meaning of *information processing* in the computational theory of mind. What kind of information is processed? The information in question need not be semantic, as not all symbols in computers are *about* something. The minimal notion that could suffice for our purposes is one of structural information: a vehicle can bear structural information in the event that it has at least one degree of freedom, that is, it may vary its state (MacKay, 1969). The number of degrees of freedom, or yes-no questions required to exactly describe its current state, is the amount of structural information. As long as there are vehicles with multiple degrees of freedom and they are part of causal processes that cause some other vehicles—just like some models of computation describe these processes

(Miłkowski, 2014)—there is information processing. This is a very broad notion, as all physical causation implies information transfer and processing in this sense (Collier, 1999).

The Gibsonian notion of information pickup requires vehicles of structural information as well. There needs to be some information out there to be picked up, and organisms have to be structured so as to be able to change their state in response to information. Gibsonians could, however, claim that the information is not processed. It is unclear what is meant by this: for example, Chemero seems to imply that processing amounts to adding more and more layers of information, like in Marr's account of vision (Chemero, 2003, p. 584; cf. Marr, 1982). But information processing need not require multiple stages of adding more information. To sum up: the Gibsonian account does not invalidate computationalism at all.

Consciousness is not computational

Some find (some kinds of) consciousness to be utterly incompatible with computationalism, or at least, unexplainable in purely computational terms (Chalmers, 1996). The argument is probably due to Leibniz's thought experiment in *Monadology* (Leibniz, 1991). Imagine a brain as huge as a mill, and enter it. Nowhere in the interplay of gears could you find perceptions, or qualitative consciousness. Hence, you cannot explain perception mechanically. Of course, this Leibnizian argument appeals only to some physical features of mechanisms, but some still seem to think that causation has nothing to do with qualitative consciousness.

The argument, if cogent, is applicable more broadly, not just to computationalism; it is supposed to defeat reductive physicalism or materialism. For this reason, this objection might be dismissed as attacking any scientific project that explains consciousness reductively.

Virtually all current theories of consciousness are computational, even the ones that appeal to quantum processes (Hameroff, 2007). For example, Bernard Baars offers a computational account in terms of the global workspace theory (Baars, 1988; cf. also Dennett, 2005), David Rosenthal gives an account in terms of higher-level states (cf. Cleeremans, 2005; Rosenthal, 2005), and Giulio Tononi explains in terms of minimal information integration (Tononi, 2004). Is there any theory of consciousness that is not already computational?

John Searle, however, suggests that only a non-computational theory of consciousness can succeed. His claim is that consciousness is utterly biological (Searle, 1992). How does this contradict computationalism given that there might be biological computers? Moreover, Searle fails to identify the specific biological powers of brains that make them conscious. He just passes the buck to neuroscience, which often offers computational accounts.

Computer models ignore time

Proponents of dynamical accounts of cognition stress that Turing machines do not operate in real time. This means that this classical model of computation does not appeal to real time; instead, it operates with the abstract notion of a computation step. There is no continuous time flow, just discrete clock ticks in a Turing Machine (Bickhard & Terveen, 1995; Wheeler, 2005). This is true. But is this an objection against computationalism?

First, some models of computation appeal to real time (Nagy & Akl, 2011), so one could use such a formalism. Second, the objection seems to confuse the formal model of computation with its physical realization. Physical computers operate in real time, and not all models of computation are made equal; some will be relevant to the explanation of cognition, and some may only be useful for computability theory. A mechanistically-adequate model of computation that describes all relevant causal processes in the mechanism is required for explanatory purposes (Miłkowski, 2014).

Brains are not digital computers

Universal Turing machines are crucial to computability theory. One could, however, maintain that brains are not digital computers (Edelman, 1992; Lupyan, 2013).

But computationalism can appeal to models of analog computation (e.g., Siegelmann, 1994), or even more complex kinds of computation (Piccinini & Bahar, 2013), if required. These models are still understood as computational in computability theory, and some theorists indeed claim that the brain is an analog computer, which is supposed to allow them to compute Turing-incomputable functions. Thus, one cannot dismiss all kinds of computationalism by saying that the brain is not a digital computer. There are analog computers, and an early model of a neural network, Perceptron, was analog (Rosenblatt, 1958). The contention that computers have to be digital is just dogmatic.

Genuine artificial intelligence is impossible

There are a number of arguments of a form:

People ψ.
Computers will never ψ.

So, artificial intelligence is impossible (or computationalism is false).

This argument is enthymematic, but the conclusion follows with a third assumption: if artificial intelligence is possible, then computers will ψ. The plausibility of the argument varies from case to case, depending on what you fill for ψ. For years, it was argued that winning in chess is ψ (Dreyfus, 1979), but it turned out to be false. So, unless there is a formal proof, it's difficult to treat premise 2 seriously.

What could be plausibly substituted for *ψ?* There are many properties of biological organisms that simply seem irrelevant to this argument, including exactly the same energy consumption, having proper names, spatiotemporal location, *etc*. The plausible candidate for substitution is some capacity for information-processing. If there is such a human capacity that computers do not possess, then the argument is indeed cogent.

Only people can see the truth A classical anti-computational argument points to the human ability to recognize the truth of logical statements that cannot be proven by a computer (Lucas, 1961; Penrose, 1989). It is based on the alleged ability of human beings to understand that some statements are true, which is purportedly impossible for machines (this argument is based on the Gödel proof of incompleteness of the first-order predicate calculus with basic arithmetic). The problem is that this human understanding has to be non-contradictory and certain. But Gödel has shown that in general it cannot be decided whether a given system is contradictory or not. So either it's mathematically certain that human understanding of mathematics is non-contradictory, which makes the argument inconsistent as it cannot be mathematically certain because it's undecidable; or the argument just assumes non-contradiction of human understanding, which makes the argument unsound because people make contradictions unknowingly (Krajewski, 2007; Putnam, 1960).

Common sense cannot be formalized Another similar argument points to common sense, which is a particularly difficult capacity. The trouble with implementing common sense on machines is sometimes called (somewhat misleadingly, cf. (Shanahan, 1997)) *the frame problem* (Dreyfus, 1972, 1979; Wheeler, 2005). Inferential capacities of standard AI programs do not seem to follow the practices known to humans, and that was supposed to hinder progress in such fields as high-quality machine translation (Bar-Hillel, 1964), speech recognition (held to be immoral to fund (Weizenbaum, 1976)), and so on. Even if IBM Watson wins in Jeopardy!, one may still think it's not enough. Admittedly, common sense is a plausible candidate in this argument.

Even if the proponent of computationalism need not require that genuine AI be based on a computer simulation of human cognitive processes, he or she still must show that human common sense can be simulated on a computer. Whether it can or not is still a matter of debate.

Computers are everywhere

At least some plausible theories of physical implementation of computation lead to the conclusion that all physical entities are computational (this stance is called *pancomputationalism,* (cf. Müller, 2009)). If this is the case, then the computational theory of mind is indeed trivial, as not only brains are computational, but also cows, black holes, cheese sandwiches etc. are all computers. However, a pancomputationalist may reply by saying that there are different kinds (and levels) of computation, and brains do not execute all kinds of computation at the same time (Miłkowski, 2007). So not just any computation but some *non-trivial kind* of computation is specific to brains. Only the kind of pancomputationalism that assumes that everything computes all kinds of functions at the same time is catastrophic, as it makes physical computation indeed trivial (Putnam, 1991; Searle, 1992).

There are no computers

Another more radical move is to say that computers do not really exist; they are just in the eyes of beholder. According to John Searle, the beholder decides whether a given physical system is computational, and therefore may make this decision for virtually everything. Nothing intrinsically is a computer. But the body of work on physical computation in the last decade or so has been focused on showing why Putnam and Searle were wrong in some sense (Chalmers, 2011; Chrisley, 1994; Copeland, 1996; Miłkowski, 2013; Piccinini, 2015; Scheutz, 1996; Shagrir, 2010b). The contemporary consensus is that computational models can be used to adequately describe causal connections in physical systems, and that these models can also be *falsely* ascribed. In other words, computational models are not different in kind from any mathematical model used in science. If they are mere subjective metaphors and don't describe reality, then mathematical models in physics are subjective as well (McDermott, 2001).

Intuitively, arguments presented by Searle and Putnam are wrong for a very simple reason: why buy a new computer instead of ascribing new software to the old one? We know that such ascriptions would be extremely cumbersome. Therefore, there must be a flaw in such arguments, and even if the technicalities involved are indeed interesting, they fail to establish a conclusion.

Conclusion

In this paper, I have listed and summarized a number of arguments against computationalism. The only objection that seems to be plausible at first glance is the one stating that common sense is impossible or extremely difficult to implement on a machine. However, more and more commonsensical capacities are being implemented on machines.

The point is that there's no good reason to think that the brain is not a computer. But it isn't a *mere* computer: It is physically embedded in its environment and interacts physically with its body, and for that, it also needs a peripheral nervous system (Aranyosi, 2013) and cognitive representations. Yet there's nothing that denies computationalism here. Most criticisms of computationalism therefore fail, and sticking to them is probably a matter of ideology rather than rational debate.

Acknowledgments

References

Aranyosi, I. (2013). *The peripheral mind: philosophy of mind and the peripheral nervous system*. New York, NY: Oxford University Press.

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge / New York: Cambridge University Press.

Bar-Hillel, Y. (1964). A Demonstration of the Nonfeasibility of Fully Automatic High Quality Translation. In *Language and Information* (pp. 174–179). Reading, Mass.: Addison-Wesley.

Barrett, L. (2015). Why Brains Are Not Computers, Why Behaviorism Is Not Satanism, and Why Dolphins Are Not Aquatic Apes. *The Behavior Analyst*, 1–15. https://doi.org/10.1007/s40614-015-0047-0

Barrett, L., Pollet, T. V., & Stulp, G. (2014). From computers to cultivation: reconceptualizing evolutionary psychology. *Frontiers in Psychology*, *5*, 867–867. https://doi.org/10.3389/fpsyg.2014.00867

Bickhard, M. H. (2008). The interactivist model. *Synthese*, *166*(3), 547–591. https://doi.org/10.1007/s11229-008-9375-x

Bickhard, M. H., & Terveen, L. (1995). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. North-Holland.

Block, N. (1995). The mind as the software of the brain. In D. Osherson, L. Gleitman, & S. Kosslyn (Eds.), *An Invitation to Cognitive Science*. Cambridge, Mass.: MIT Press.

Chalmers, D. J. (1996). *The conscious mind: in search of a fundamental theory*. New York: Oxford University Press.

Chalmers, D. J. (2011). A Computational Foundation for the Study of Cognition. *Journal of Cognitive Science*, (12), 325–359.

Chemero, A. (2003). Information for perception and information processing. *Minds and Machines*, *13*, 577–588.

Chrisley, R. L. (1994). Why everything doesn't realize every computation. *Minds and Machines*, *4*(4), 403–420. https://doi.org/10.1007/BF00974167

Cleeremans, A. (2005). Computational correlates of consciousness. *Progress in Brain Research*, *150*, 81–98. https://doi.org/10.1016/S0079-6123(05)50007-4

Collier, J. D. (1999). Causation is the transfer of information. In H. Sankey (Ed.), *Causation, natural laws and explanation* (pp. 279–331). Dordrecht: Kluwer.

Copeland, B. J. (1996). What is computation? *Synthese*, *108*(3), 335–359.

Cummins, R., & Roth, M. (2012). Meaning and Content in Cognitive Science. In R. Schantz (Ed.), *Prospects for Meaning* (pp. 365–382). Berlin & New York: de Gruyter.

Daugman, J. (1990). Brain metaphor and brain theory. In E. L. Schwartz (Ed.), *Computational Neuroscience* (pp. 9–18). Cambridge, Mass: MIT Press.

Dennett, D. C. (2005). *Sweet Dreams. Philosophical Obstacles to a Science of Consciousness*. Cambridge, Mass.: MIT Press.

Dretske, F. I. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: form, content, and function* (pp. 17–37). Oxford: Clarendon Press.

Dreyfus, H. (1972). *What Computers Can't Do: A Critique of Artificial Reason*. New York: Harper & Row, Publishers.

Dreyfus, H. (1979). *What computers still can't do: a critique of artificial reason*. Cambridge Mass.: MIT Press.

Edelman, G. M. (1992). *Bright air, brilliant fire: on the matter of the mind*. New York, N.Y.: BasicBooks.

Egan, F. (1995). Computation and Content. *The Philosophical Review*, *104*(2), 181–181. https://doi.org/10.2307/2185977

Ekman, P. (2003). *Emotions revealed: recognizing faces and feelings to improve communication and emotional life*. New York: Times Books.

Epstein, R. (2016, May 18). The empty brain. Retrieved December 28, 2016, from https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer

Fresco, N. (2010). Explaining Computation Without Semantics: Keeping it Simple. *Minds and Machines*, *20*(2), 165–181. https://doi.org/10.1007/s11023-010-9199-6

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

Fresco, N., & Wolf, M. J. (2013). The instructional information processing account of digital computation. *Synthese*, *191*(7), 1469–1492. https://doi.org/10.1007/s11229-013-0338-5

Gallistel, C. R., & King, A. P. (2010). *Memory and the Computational Brain*. Chichester: Wiley-Blackwell.

Gibson, J. J. (1986). *The Ecological Approach to Visual Perception*. Hove: Psychology Press.

Hameroff, S. R. (2007). The Brain Is Both Neurocomputer and Quantum Computer. *Cognitive Science*, *31*, 1035–1045.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, *42*, 335–346.

Haugeland, J. (1985). *Artificial intelligence: the very idea*. Cambridge, Mass.: MIT Press.

Karl, F. (2012). A Free Energy Principle for Biological Systems. *Entropy (Basel, Switzerland)*, *14*(11), 2100–2121. https://doi.org/10.3390/e14112100

Krajewski, S. (2007). On Gödel's Theorem and Mechanism: Inconsistency or Unsoundness is Unavoidable in any Attempt to "Out-Gödel" the Mechanist. *Fundamenta Informaticae*, *81*(1), 173–181.

Lakoff, G. (1987). *Women, fire, and dangerous things: what categories reveal about the mind*. Chicago: University of Chicago Press.

Leibniz, G. W. (1991). *The monadology*. (R. Latta, Trans.). Raleigh, N.C.; Boulder, Colo.: Alex Catalogue; NetLibrary.

Lucas, J. (1961). Minds, Machines and Gödel. *Philosophy*, *9*(3), 219–227.

Lupyan, G. (2013). The difficulties of executing simple algorithms: Why brains make mistakes computers don't. *Cognition*, *129*(3), 615–36. https://doi.org/10.1016/j.cognition.2013.08.015

MacKay, D. M. (1969). *Information, mechanism and meaning*. Cambridge: M.I.T. Press.

Marr, D. (1982). *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman and Company.

McDermott, D. V. (2001). *Mind and Mechanism*. Cambridge, Mass.: MIT Press.

Miłkowski, M. (2007). Is computationalism trivial? In G. D. Crnkovic & S. Stuart (Eds.), *Computation, Information, Cognition – The Nexus and the Liminal* (pp. 236–246). Newcastle: Cambridge Scholars Press.

Miłkowski, M. (2013). *Explaining the Computational Mind*. Cambridge, Mass.: MIT Press.

Miłkowski, M. (2014). Computational Mechanisms and Models of Computation. *Philosophia Scientae*, *18*(18–3), 215–228. https://doi.org/10.4000/philosophiascientiae.1019

Millikan, R. G. (1984). *Language, thought, and other biological categories: new foundations for realism*. Cambridge, Mass.: The MIT Press.

Müller, V. C. (2009). Pancomputationalism: Theory or metaphor? In R. Hagengruber (Ed.), *The relevance of philosophy for information science*. Berlin: Springer.

Nagy, N., & Akl, S. (2011). Computations with Uncertain Time Constraints: Effects on Parallelism and Universality. In C. Calude, J. Kari, I. Petre, & G. Rozenberg (Eds.), *Unconventional Computation* (Vol. 6714, pp. 152–163). Springer Berlin / Heidelberg. Retrieved from http://dx.doi.org/10.1007/978-3-642-21341-0_19

Newell, A. (1980). Physical symbol systems. *Cognitive Science: A Multidisciplinary Journal*, *4*(2), 135–183. https://doi.org/10.1207/s15516709cog0402_2

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Penrose, R. (1989). *The emperor's new mind*. London: Oxford University Press.

Piccinini, G. (2008). Computation without Representation. *Philosophical Studies*, *137*(2), 205–241. https://doi.org/10.1007/s11098-005-5385-4

Piccinini, G. (2010). The Mind as Neural Software? Understanding Functionalism, Computationalism, and Computational Functionalism. *Philosophy and Phenomenological Research*, *81*(2), 269–311. https://doi.org/10.1111/j.1933-1592.2010.00356.x

Piccinini, G. (2015). *Physical computation: a mechanistic account*. Oxford: Oxford University Press.

Piccinini, G., & Bahar, S. (2013). Neural computation and the computational theory of cognition. *Cognitive Science*, *37*(3), 453–88. https://doi.org/10.1111/cogs.12012

Preston, J., & Bishop, M. (2002). *Views into the Chinese room: new essays on Searle and artificial intelligence*. Oxford; New York: Clarendon Press.

Putnam, H. (1960). Minds and machines. In S. Hook (Ed.), *Dimensions of Mind*. New York University Press.

Putnam, H. (1991). *Representation and Reality*. Cambridge, Mass.: The MIT Press.

Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, Mass.: MIT Press.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, *65*(6), 386–408. https://doi.org/10.1037/h0042519

Rosenthal, D. (2005). *Consciousness and mind*. Oxford, New York: Oxford University Press.

Scheutz, M. (1996). When Physical Systems Realize Functions…. *Minds and Machines*, *9*(2), 1–34. https://doi.org/10.1023/A:1008364332419

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 1–19. https://doi.org/10.1017/S0140525X00005756

Searle, J. R. (1990). Is the Brain's Mind a Computer Program? *Scientific American*, (January), 26–31.

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

Shagrir, O. (2010a). Brains as analog-model computers. *Studies In History and Philosophy of Science Part A*, *41*(3), 271–279. https://doi.org/10.1016/j.shpsa.2010.07.007

Shagrir, O. (2010b). Towards a Modeling View of Computing. In G. Dodig-Crnkovic & M. Burgin (Eds.), *Information and Computation*. Singapore: World Scientific Publishing.

Shanahan, M. (1997). *Solving the frame problem: a mathematical investigation of the common sense law of inertia*. Cambridge, Mass.: MIT Press.

Siegelmann, H. (1994). Analog computation via neural networks. *Theoretical Computer Science*, *131*(2), 331–360. https://doi.org/10.1016/0304-3975(94)90178-3

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, *5*(1). https://doi.org/10.1186/1471-2202-5-42

Weizenbaum, J. (1976). *Computer power and human reason: from judgment to calculation*. San Francisco: W.H. Freeman.

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

Zauner, K.-P., & Conrad, M. (1996). Parallel computing with DNA: Toward the anti-universal machine. In H.-M. Voigt, W. Ebeling, I. Rechenberg, & H.-P. Schwefel (Eds.), *Parallel Problem Solving from Nature - PPSN IV* (Vol. 1141, pp. 696–705). Springer Berlin / Heidelberg.