Philosophy of AI
Vol. 2, 2026, 43-60
DOI: 10.18716/ojs/phai/2026.11564

NEUROMORPHIC COMPUTING AND PRACTICAL MEDIUM DEPENDENCE

Mazviita Chirimuuta
University of Edinburgh, Department of Philosophy, Edinburgh, United Kingdom

Abstract
The increasingly prohibitive cost of energy demanded by large artificial neural networks (ANNs) is giving new impetus to research and development on neuromorphic computing. Importantly, there is an open question over how brain-like the hardware will have to be in order for an artificial intelligence to match the brain in its combination of robustness, adaptability, and energy efficiency. If biological cognition is heavily dependent on the specific properties of the material that instantiates it (i.e. living cells), then neuromorphic computing will have to merge with synthetic biology in order to achieve its ultimate goal of brain-like performance. If it is not, neuromorphic computing holds out the promise of some gains in efficiency but there is no pressure for hardware to become increasingly neuro-mimetic in order to match the functionality of the nervous system. In this paper I introduce the concept of *practical medium dependence/independence* in order to explore the likelihood of these two scenarios. I present the argument that practically medium independent approaches to information processing, such as digital computing, are inherently less efficient than ones dependent on the specifics of implementing media, and for that reason will not have evolved. This result has implications for how we rate the near-term possibility of human-like artificial general intelligence.

Keywords: Philosophy of Neuroscience; AI; Medium Independence; Evolution of Cognition; Neuromorphic Computing

1. Introduction

The increasingly prohibitive cost of energy demanded by large artificial neural networks (ANNs) is giving new impetus to research and development on neuromorphic computing. The immensity of the difference in energy requirements of a human brain, versus a large language model, is often noted and I will not rehearse all the figures here.[1] The field of neuromorphic engineering encompasses an array of designs that are somewhat inspired by certain features of neural systems. Importantly, there is an open question over how brain-like the hardware will have to be in order for an artificial intelligence to match the brain in its

---

[1] See e.g. Jamadar et al. (2025, 7): "the brain nonetheless runs continuously on only ~17 watts of power. By comparison, a large high-performance computing cluster uses up to six orders of magnitude more power, operating at ~2 megawatts."

combination of robustness, adaptability, and energy efficiency. If, on the one hand, biological cognition is heavily dependent on the specific properties of the material that instantiates it (i.e. living cells), then neuromorphic computing will have to merge with synthetic biology in order to achieve its ultimate goal of brain-like performance. If, on the other hand, biological cognition is relatively independent of its medium then neuromorphic computing holds out the promise of gains in efficiency but there is no pressure for hardware to become increasingly neuro-mimetic in order to match the functionality of the nervous system. Philosophy can play a role here in evaluating the likelihood of these outcomes.[2] There are ongoing debates over whether neural computing is medium dependent or independent, and whether computing is essentially medium independent.[3] The purpose of this paper is to make an intervention in the first of these two disputes, focussing us on the pertinent issue of feasibility by introducing the new concept of *practical medium dependence/independence.* This concept tracks the dependence or independence of the functionality and performance of an information processing system on the specific details of its physical realiser. With this notion in hand I will be arguing that the former scenario, in which neuro-mimesis is required for the achievement of brain-like functionality, is the more likely one.

Accordingly, the first task of the paper will be to discuss existing notions of medium independence and why the concept of practical medium independence is required. The basic point is that it is beneficial, when considering the variety and advantages of neuromorphic and neuronal media, to shift attention away from ‘in principle’ issues around computability and towards ‘in practice’ issues about the feasibility of implementing certain processes in different media. A system is practically medium independent if there is a formal description of the system that can feasibly be implemented in a fundamentally different material substrate, and that description is rich enough to account for the system’s observed behaviour. This concept will be the subject of Section 2.

In Section 3 I present an argument that brain-based cognition is not practically medium independent because the costs of practically medium independent computation are prohibitive in comparison to the medium dependent solutions that were available to evolution. Medium dependent processes exploit the specific properties of the implementing medium in order to achieve efficient use of material and energetic resources. That is why they cannot be implemented in alternative media without drastic increases in costs. In Section 4 I discuss the implications of this result for the engineering fields of neuromorphic computing and artificial general intelligence (AGI). Since it suggests that brain-like functionality can only be achieved with brain-like substrate (i.e. living neuronal tissue) we should therefore expect that human-like AGI will not happen unless there are dramatic breakthroughs in bioengineering or novel inorganic computing methods. In the remainder of this introduction I provide some more background on neuromorphic computation and relate our topic to the theory of computational complexity.

1.1 Costs of Computing

The costs of computing are high. Back in 2016 the world’s data centres consumed about 40% more energy than the entire United Kingdom.[4] In 2022, 5-15% of the world’s energy budget was estimated to be spent on data manipulation (Christensen et al. 2022, 3). This is an ever expanding figure due, in particular, to the energetic demands of AI based on large language

---

[2] Granted, there is a spectrum of intermediate possibilities. For ease of exposition I will talk as if this were a clear dichotomy.

[3] See Maley (2021, 2025), Piccinini (2015, 2020), Anderson & Piccinini (2024), Drayson (2025).

[4] https://www.rankred.com/largest-data-centers-in-the-world/

models (LLMs). In the context of scaling up of LLMs being pursued as the most direct path towards AGI, there is a reasonable fear that progress will be inhibited unless there is a breakthrough towards utilising new energy sources such as cold fusion.[5] The human brain is a costly organ – it is estimated to account for 20% of the body's resting metabolism (Jamadar et al. 2025, 1) – but it is nowhere near as energetically demanding as a supercomputer. For this reason, the research field of neuromorphic computing has been in large part motivated by the promise of greater energy efficiency through both hardware and software design that takes direct lessons from biology (Schuman et al., 2017).

The term neuromorphic computing encompasses a large range of digital, analogue, and hybrid approaches that depart from the classic von Neumann architecture (Schuman et al., 2022). A commonality that will be crucial to the argument of this paper is that all such designs are brain-like only insofar as we abstract away from chemical signalling in the brain and identify neural processing with electrical impulses (i.e. spikes). Spiking artificial neural networks are a popular neuromorphic approach. In terms of computer substrate or implementing media, most use conventional silicon and metal oxide semiconductor technology (Schuman et al 2022:11). Failure to achieve brain-like energy efficiency using conventional hardware has motivated research into novel substrates, using organic (but non-living) and inorganic materials that attempt to model neuronal features such as ion channels (Najem et al., 2018, 4702-4703).

Energy costs are a practical concern and philosophers of mind have tended to be somewhat dismissive of these kinds of constraint, subject, as they are, to be loosened as technology advances. The results of computability theory, such as the Turing-Church thesis, are foundational both to Putnam's original machine functionalism and to those who argue today that human-like AGI can be implemented in silicon hardware that is materially unlike neural tissue, or even that it is possible to upload a human mind into the cloud.[6] Here computation, qua mathematical process, is defined entirely independently of any material that implements it. As we will discuss in the next section, this is the basis for the view that computation (and by extension cognition) is essentially medium independent.

The theory of computational complexity is another branch of mathematics, one that does bear on practical considerations. Even if it can be shown in principle that an algorithm stops in a finite number of steps, the algorithm is of no use if the time to completion will be greater than the age of the universe. As reviewed by Aaronson (2013), complexity theory determines the efficiency of algorithms in terms of scaling of resources need to solve a problem with a measure of the problem size. Inefficient algorithms that would quickly use up resources available for computation (time, energy and space available to humans on a finite planet) are classified *unfeasible*. One reason why considerations of feasibility have been paid less attention in philosophy of mind than those of computability is that 'in principle' grounds for the computational theory of mind have been presumed more philosophically significant than 'merely' practical issues.[7]

---

5. Currently, nuclear reactors are being recommissioned in order to increase electricity supply to data centres (Greshko, 2022).

6.(Chalmers, 2014) and see (Pigliucci, 2014) for a biological counter-argument, congenial to the view presented in this paper.

7. Aaronson (2013:265) makes the important observation that questions of feasibility tend to be more scientifically relevant than questions of computability:
"insisting that programs terminate after reasonable amounts of time, that they use reasonable amounts of memory, etc. might sound like relatively-minor amendments to Turing's notion of computation. In practice, though, these requirements lead to a theory with a completely different character than computability theory. Firstly, complexity has much closer connections with the sciences: it lets us pose questions about (for example) evolution, quantum mechanics, statistical physics, economics, or human language acquisition that would be meaningless from a computability standpoint (since all the relevant problems are computable)."

Some computer scientists and cognitive scientists have begun to apply complexity theory to the question of AGI (Rooij et al., 2024). Here I will not be presenting mathematical results in this manner, drawing instead on biological principles and empirical results from neuroscience. [8] The approach is prefigured by that of Thagard (2022) who makes the case for the importance of energy constraints to evaluation of the functionalist thesis that human cognition is "substrate independent": even if human cognition is a kind of computation and therefore theoretically independent its substrate or medium, it may not be feasible to implement those same computations in non-neuronal materials because of the energy demanded by the alternative implementations.[9] The notion of practical medium dependence, presented in the next section, is a way to pin down this kind of constraint.

Computational unfeasibility equates with nomological impossibility when it can be shown that completion of an algorithm would demand cosmic scales of time, energy and space – e.g. a computer the size of a planet, using all the sun's energy for a billion years. However, I will argue in the concluding section that even more modest results are still interesting because they can reveal much about the difference between biological and engineered methods of information processing. Of course, there is a potential instability in judgments of infeasibility of a computational system where economic and material constraints currently, and in the foreseeable future, put the matter outside the scope of human possibility, but where one might envisage a technological breakthrough (either in power generation or computing machinery) that could lift those constraints. In my view, it is unhelpful to reach for science fiction scenarios at this juncture: they are more likely to obfuscate important lessons to be taken from biology and neuroscience than to reveal general principles about cognition and its material basis.

2. Medium Independence and Practical Medium Independence

The term *medium independence* was introduced to the philosophy of computing by John Haugeland as part of a characterization of computers as *automatic formal systems*:

> Formal systems are independent of the medium in which they are 'embodied'. In other words, essentially the same formal system can be materialized in any number of different media, with no formally significant difference whatsoever. This is an important feature of formal systems in general. I call it *medium independence.* (Haugeland, 1985, 58)

On Haugeland's account, computers are physical systems that implement formal systems and because of medium independence, the only theoretical restriction on the physical system is that it has sufficient number of degrees of freedom to instantiate the tokens of the formal system (such as zeros and ones in a binary digital computer) that are arbitrarily assigned to its components. When characterizing the functionality of a concrete system that implements a formal system, it is therefore legitimate to abstract away from all of its specific material properties and refer only to structural features that track relationships between the tokens. This relates to the familiar idea that we characterize the functionality of a computer

---

[8] I should emphasise an important difference between my approach and computational complexity analysis. Those analyses indicate the resources demanded by an algorithm in *any medium*. The actual choice of medium will determine whether a less efficient algorithm can still be run – the faster the hardware, the more feasible it will be. As with complexity analysis, my interest is in feasibility. But, specifically, my focus is on the properties of biological 'hardware' and the hypothesis that it makes feasible certain processes that cannot be implemented in alternative media without soon outstripping available resources. See also van Rooij (2008)

[9] In this paper my intention is to be neutral on the issue of whether information processing in the brain is rightly classified as computation, or should be given another term such as 'signalling'. The argument for practical medium dependence will apply in either case.

with the abstract description given in the software, not with a detailed description of the hardware. (Haugeland, 1997, 10–11) therefore writes that,

> A *concrete system* [emphasis added] is medium independent if what it is does not depend on what physical "medium" it is made of or implemented in. Of course, it has to be implemented in *something;* and, moreover, that something has to support whatever structure or form is necessary for the kind of system in question. But, apart from this generic prerequisite, nothing specific about the medium matters (except, perhaps, for extraneous reasons of convenience).

This stems directly from the conception of computing machines as physical systems that implement formalisms such as Turing machines and finite state automata. These formalisms are mathematical abstracta and as such are essentially unrelated to their physical instantiations. All that matters, in principle, about the physical medium is that it has a rich enough causal structure to map on to the structure of the computation that it implements, often only approximately. In practice, of course, there are many constraints on which materials can be used to build usable computing systems but these get put under the heading of "extraneous reasons of convenience".

As discussed by (Williams, 2025), Haugeland's notion of MI differs from that more recently employed by Gualtiero Piccinini (2015, 2020) within the mechanistic account of physical computation. Piccinini (2015, 2020) has argued that the manipulation of medium independent vehicles is an essential feature of physical computation and has contrasted medium independence with the weaker constraint of multiple realizability. The basic difference is that a function such as trapping a mouse or taking a cork out of a bottle involves an operation on a certain kind of physical object (a mouse or a cork) and as such the mechanisms for it, while multiply realizable, are required to have certain physical properties that make the operation possible (e.g. rigidity). In contrast, medium independence puts no constraints like this on the physical properties of the implementing tokens (Anderson and Piccinini 2024, 243). In this case, the medium independence of the vehicles stems from their featuring in a highly abstract representation of a physical system (a "mechanism sketch"). Williams (2025) calls this "abstract-as-omission", contrasting it with Haugeland's MI formal systems which are "abstract-as-abstracta".

The mechanistic account of physical computation is intended to demarcate physical systems which compute (laptops and, arguably, brains) from those which do not (dinner plates and tomatoes). Although it is supposed to avoid the problem of positing an implementation relation – a problem which arises when physical computers are conceptualised as concrete systems that realize abstract formal systems – it is unclear that it can actually afford to discard the notion of implementation (Kersten, 2024; Kuokkanen, 2022). Hence various authors, such as Drayson (2025), still assert that a claim that a physical system such as a brain performs computations requires a claim that it implements a formal system.

2.1 Medium Independence and the Brain

Piccinini has argued that because cognitive processes are computational, they are also medium independent. In support of this he asserts that neuronal spikes are medium independent (Piccinini 2020, 213-214) and argues that the success of deep convolutional network models of the visual system also demonstrates the medium independence of neural computation. As Anderson and Piccinini (2024, 254) write,

> neural computation can be studied largely in abstraction from the concrete features of the physical medium that implements it, so that the same or similar enough computations can be implemented in artificial neural networks and the similarities between biological and artificial neural computations can be used to shed light on both. This illustrates the medium independence of neural computation.

Cao (2022) and Chirimuuta (2022) present an empirical case against the medium independence of neural processing because of the implication, apparent in the passage just quoted, of the irrelevance of the material features of neuronal tissue to the investigation of the actual functions undertaken by the brain. A particular focus is the role of neurochemistry for signalling within the brain. The fact that voltages can be generated through various materials lends credence to the idea that electrical signalling is medium independent. However, spiking is utterly enmeshed with chemical signalling, the most obvious point being that the end result of spiking is transmitter release at the synapse.[10] Chemical signalling is obviously dependent on the specific features of the molecules involved, such as the shape of particular proteins and the binding affinities of amino acids; the idea that it could be medium independent has no traction.

However, Drayson (2025) has argued that this line of argument rests on a misunderstanding of the claim of medium independence: of course the physical properties of neural tissue make a difference to how information processing occurs – the property of electrical conductivity being an obvious instance – but that does not stand against the point that computation is the implementation of a formal system and as such is not inherently dependent on any physical properties, so that computation in neural tissue is just as much medium independent as computation in electronics.[11]

Still, this prompts further reflection over why we should think that the brain implements any formal system. Abstract Turing Machines preceded the invention of actual digital computers which instantiate them only approximately (since none have infinite memory). Digital computers were invented specifically to implement formal systems; evolution did not have formal systems in mind, prompting the production of brains for their implementation. As argued by Maley (2023), computational descriptions of the brain are models devised by neuroscientists based on recorded data and theoretical hypotheses. They are abstractions from actual neural mechanisms but the existence of an abstraction in this sense of *omission* of physical details does not imply the existence of an abstraction in the *mathematical* sense of a formal system. Unfortunately, many have ignored the difference between these two senses of abstraction and adhere to the view that the brain implements formal systems which the best neurocomputational models approximately capture. Since formal systems are medium independent by definition it follows that *in principle* any of these putative neuro-computations could be implemented in a physical system utterly unlike the brain and the question of practical feasibility is normally ignored by philosophers, though not for working neuromodellers who have to budget run time on computers for large and detailed brain models.

In this context, the assertion of medium independence of neural processes becomes an empty claim. It rests on the conflation of abstractions by omission with mathematical abstractions (Polger & Shapiro, 2023, 332-333). By definition, abstract formal systems are

---

[10] In fact, the synapses which involve the axon terminal of a spiking neuron are referred to as *chemical* synapses, whereas *electrical* synapses are gap junctions, where no action potentials are involved (Pareda, 2014).

[11] And see Piccinini (2022). Note, also, that Chirimuuta (2022) explictly rejects the assertion made by Drayson, that computation is essentially medium independent, siding with Maley's (2021) view that analogue computation is medium dependent and arguing that neural computation would be more similar to that case; her paper is therefore not guilty of confusion about the claim for medium independence, but simply disagrees with Drayson's definition.

medium independent but we have not been given evidence that these have anything to do with the operation of the brain – this evidence is not supplied by the existence of computational models of neural systems, which are abstract models in the sense that they are constructed by omitting most of the details of systems (Chirimuuta, 2021, 2024, chap. 4). Furthermore, the claim is empty because even if we grant that brains implement MI formal systems, there is the further question of whether such formal systems could in practice be implemented in any other medium. The question of feasibility needs to be addressed. Even if the brain does house a formal system, alternative implementations are not on the horizon if they would take a data centre the size of Asia.

2.2 Practical Medium Independence

As an alternative to this unfruitful framing of the issues, I introduce the notion of *practical medium independence (prac-MI)*. The functionality[12] of a concrete system is prac-MI iff:

1. There is a formal description of the system which can feasibly be implemented in a fundamentally different material substrate,
2. and this description is rich enough that it accounts for all the functional capacities of the system that have been empirically ascertained.

By fundamentally different material substrate I mean the difference between inorganic chemistry and carbon based organic chemistry, or between these and the biochemistry of living cells. By feasible I mean nomologically possible within a reasonable estimate of resources available to humans, bracketing speculation about what would become possible with unrealistic new technologies. Given that (1) is about a mathematical abstraction, this formal description could be as complex as you like – it could be isomorphic with a structure that maps every molecule and interaction in the concrete system. The catch is that the more complex it is, the less likely it is that it can feasibly be implemented in a fundamentally different substrate. I am here ignoring trivial implementations like Putnam's rock and Searle's brick wall (Sprevak, 2018) – implementation of the formal system is taken to require a mapping to the actual causal structure of the concrete system (see Piccinini and Anderson 2024).[13]

Condition (2) stipulates that the formal description must be rich or complex enough to account for the observed behaviour of the concrete system. Obviously, new observations can always be made, which leaves claims for prac-MI open to empirical refutation. However, the real point of this clause is to rule out as alternative implementations ordinary *models* which are partial and known to capture only a subset of observed capacities. Scientists embrace these kinds of abstractions (in the sense of omission) for their convenience and because focussing on some but not all behaviours is legitimate in the context of research. But in our discussion what's at issue is a claim of formal equivalence between a concrete system and a candidate alternative implementation. For such claims to be made, all known capacities

---

[12] On previous definitions, MI has been used to refer to various kinds of entities: formal systems, vehicles, concrete systems (see above). I'm applying it to functionality precisely because that is what's at issue in this investigation: whether the functional capacities of a concrete system are dependent or independent on the medium. In the case of the brain, I am restricting the notion of functionality to cognitive functions. There are other biological functions that neurons perform (such as maintenance of the cell membrane) that are clearly not MI or prac-MI and we will ignore these for the time being, though in Section 4 we will see that the demarcation between cognitive and non-cognitive functions of the brain can be challenged.

[13] Because feasibility is at issue, which includes the time course of the processes, the dynamics of the concrete system matter. It could be argued that no formal description can account for these. Rather than declare that in that case prac-MI never obtains, I would rather expand the notion of the allowed formal description to encompass mathematical representation of dynamics, as commonly found in computer simulations of physical systems.

need to be accounted for, otherwise the claims risk triviality. Piccinini (2020: 42) makes a related point regarding claims for multiple realizability:

> If a description is sufficiently general, the property picked out by that description might have many lower-level realizers, but perhaps only trivially so; its multiple realizability might be an artifact of such a broad higher-level description that the property plays no useful role in scientific taxonomy or explanation.

Needless to say, medium independence in the previous senses (Williams 2025) does not imply prac-MI. Many computations that are in principle medium independent are in practice dependent on particular kinds of hardware for feasible implementation: you couldn't build a Babbage machine big and fast enough to run MS Word. That said, I take it that the functionality of existing digital and most analogue computers is prac-MI. For example, the metal disc of an analogue device for computing integrations (Maley, 2023a, 742-743) could be replaced with a rigid plastic one and the functional capacity would be unchanged. The high performance of today's digital electronics is of course dependent on careful engineering of silicon based materials. But I see no reason to think that – even if a nightmare for engineers – it would not be feasible to build machines based on different classes of chemical components that could perform the same tasks (albeit in a somewhat less efficient manner). Systems that are not prac-MI are practically medium dependent (prac-MD). Quantum computation is expected to be prac-MD, due to speed of computation impossible to achieve in classical computers. Prac-MI is evidently not a necessary condition for physical computation.[14]

Current Artificial Neural Network models of brain areas, including spiking ANNs, satisfy (1) but not (2). ANNs at best roughly track electrical signalling networks in the brain. They fail to capture too many functionally relevant details (including all the chemical ones) to be able to reproduce all but a limited number of functional capacities of the areas targeted.[15] For example, dynamics of neural plasticity depend on chemical processes that are not at all represented in these models and unsurprisingly the learning rates of these models are very different from what is observed in animals. Scientifically, omission is unproblematic; as emphasized above, all models are partial. A formal description of a brain area that mapped molecular as well as electrical signalling could potentially satisfy (2) but not (1). I refer here to Cao's (2022) impressively documented case concerning the range of cognitively significant molecular interactions in brain tissue, including those mediated by the glial cells that tend to be ignored because they are not electrically excitable. In particular, I would like to highlight Cao's point that these interactions depend on each neuron being equipped with a whole suite of chemical sensors.[16] There seems no efficient way to reproduce these sensors and interactions except by reproducing the original chemistry. Of course, simulations of biochemical processes exist for research purposes. The point is that they are so computationally intensive that they cannot feasibly be scaled up to encompass, say, the neural circuit associ-

---

[14] Arguably, deep learning AI models today are practically medium dependent because they can be feasibly implemented only on specialised hardware (Hooker, 2020). This is not a problem for my argument as it just means that dependence of functionality on substrate is a pervasive feature of high performing systems, not just biological ones.

[15] See Bowers et al. (2023) and van Rooij et al (2024) on how broad claims of functional equivalence break down if you look in details of the capacities of ANNs.

[16] This fact tends to be underplayed because of the centrality of electrophysiology in neuroscience, which is due to availability of recording technologies as well as convenience, not a reflection of any actual priority of electrical over chemical signalling in the brain. As Sala et al. (2023, 354) write:
"Because the recording of electrical signaling or its proxies has long been considered to be a reasonable strategy to investigate functional links between brain regions, chemical signaling or its proxies should represent an alternative, equally reasonable means to access patterns of neuronal communication. In fact, because chemical synapses represent by far the predominant mode of signal transduction in the human brain, neural communication may be more accurately reflected by biochemical than by electrical signaling."

ated with a particular cognitive task (Cao 2022, fn 11). Even with the extreme abstraction of the 'point neuron', the nodes used in standard ANNs running on digital graphical processing units (GPUs), scaling to brain equivalent network sizes (i.e. billions of neurons) pushes energy costs beyond the reach of all but the world's richest tech companies. The reasonable inference is that the demands of adding bio-realistic detail to ANNs at that scale would render them (nomologically) impossible to implement.

One way to object to this line of argument is to dispute the claim that these 'gory details' need to be included in a description of the brain if it is to account for all the functional capacities. There is indeed an open question within neuroscience as to how far one can abstract away from the actual molecular properties of the brain and still capture an acceptable amount of functionality.[17] The notion of prac-MD is a useful one for framing this issue, especially if we allow it to track a spectrum of weak to strong dependence of a system's functionality on its material specifics. In the next section I provide an argument that the brain is likely to be strongly prac-MD, which means that sparse abstractions are unlikely to be adequate for achieving functional equivalence. Before moving on I would like to note that the original notion of medium independence is apt for digital computing precisely because in this case a description which abstracts away from all properties of the physical instantiating system (other than degrees of freedom) still offers a complete description of the functionality. Abstraction (in the sense of omission) *aligns with* an abstraction in the mathematical sense of a formal system. This is how things were intended to be – at least that was the aspiration, though the real world history of designing AI may not have shown this (Hooker 2020). As Herbert Simon writes in his *Sciences of the Artificial*:

> No artifact devised by man is so convenient for this kind of functional description as a digital computer. It is truly protean, for almost the only ones of its properties that are detectable in its behavior (when it is operating properly!) are the organizational properties. The speed with which it performs it basic operations may allow us to infer a little about its physical components and their natural laws; speed data, for example, would allow us to rule out certain kinds of "slow" components. For the rest, almost no interesting statement that one can make about an operating computer bears any particular relation to the specific nature of the hardware. (Simon, 1969, 18)

To argue for the practical medium dependence of the brain is to insist on how different the brain is from such a machine. This is not a matter of opinion or judgment – my point will be that it would have been too costly for evolution to have operated in this way.

3. The Resource Argument for Practical Medium Dependence

As observed by Herbert Simon in the passage just quoted, digital computers were designed in such a way that their functionality could be maximally 'indifferent' to their material constitution. While instances of both analogue and digital computers can satisfy the criteria for practical medium independence, there is a difference in how these two kinds of computing relate function to constitution, and this is a helpful entry into our discussion of the link between prac-MD and resource efficiency. There are various ways to draw the analogue/dig-

---

[17] See for example (Arkhipov et al., 2018) who report surprisingly good performance of an extremely simplified cortical model, in comparison with a more detailed one.

ital distinction and these are an ongoing matter of controversy.[18] I will refer to a point made by Schuman et al (2017:12) which I take to be uncontroversial:

> Analog systems utilize native physical characteristics of electronic devices as part of the computation of the system, while digital systems tend to rely on Boolean logic-based gates, such as AND, OR, and NOT, for building computation.

The observation is that designers of analogue computers leverage specific physical properties of the machine's components in order to perform computations that are not programmable through software but embedded in the layout of the machine. These properties are *analogous* to magnitudes over which the computation is to be performed in such a way that running the machine delivers the result of a computation (Maley 2023a). In a comparison of the costs of computation for analogue and digital systems, Sarpeshkar (1998) shows how analogue can in principle use far fewer resources than digital, both in terms of energy and space (i.e. number of components). The benefit of digital is that it achieves a much better signal to noise ratio than analogue. To achieve the precision needed for most modern computing tasks, the price of digital has been worth paying. Figure 1 explains the basic intuition behind Sarpeshkar's results. Take an arbitrary piece of physical equipment like a forearm that can be moved to different angles – we will use it as a signalling system. If we devise a digital code we are restricted to using only two states of the arm to signal the presence and absence of a particular state. This equals one bit of information per arm. However, if we play around with this equipment we should notice that we are under-exploiting its signalling potential. Given that the one arm can be moved through a continuum of angles it can be employed to signal an endless number of different states. This is an analogue coding strategy. It is not just that we are we coding continuous variables but that we are leveraging specific properties of the physical medium to do so. We will get many more bits per arm, limited only by our ability to discriminate angles. In other words, we can code a lot of information with this one piece of equipment but risk a poor signal to noise ratio.

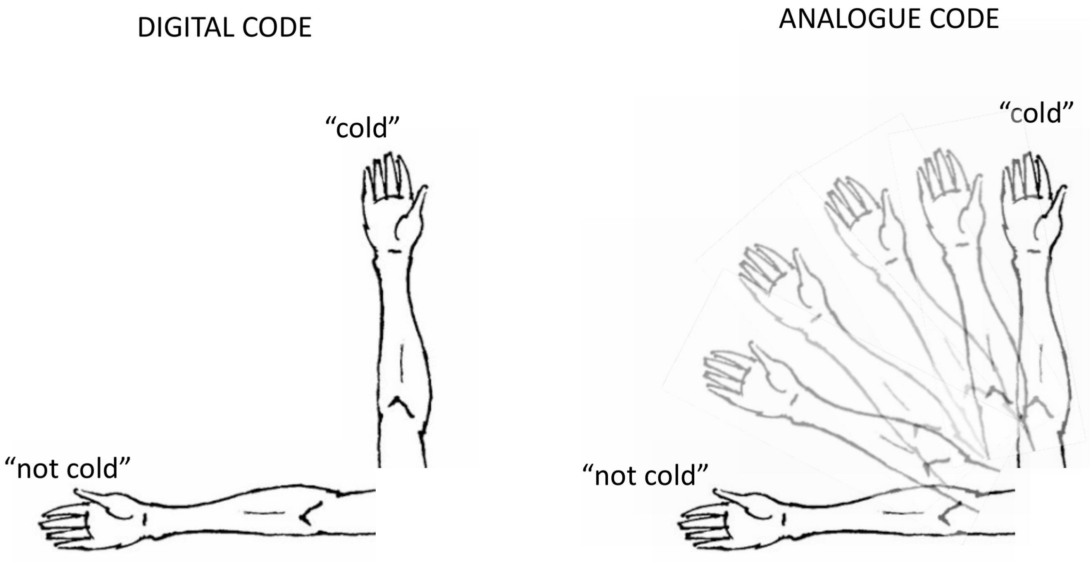

Figure 1. Figure 1: Illustration that one physical device can signal more or less information, depending on whether an analogue or digital code is employed.

---

[18] See Haugeland (1998), Kulvicki (2016), and Maley (2011).

This is all by way of preamble to a discussion of signalling in the brain. As mentioned before, both analogue and digital computers can be practically medium independent (and I am not arguing that the brain is an analogue computer), but we can still note that between these two classes of computers there is a difference in how efficiently they exploit their material resources. One wants to say that the analogue computation is in some sense more dependent on its medium, and this is the source of its greater efficiency. I began with a toy example to illustrate a difference in principle. Examples could be multiplied – take any one given physical system – a candle, a shoe or a flag – and ask yourself how you might use it as the medium for a code. If you go digital, you force the thing to switch between two states; if you go analogue, you play around with its native capacities, including its dynamics, and see how these can be used to code a far greater number of states. Hence with one and the same physical resource, more or less information can be coded. The crucial point is that you increase informational capacity by exploiting the peculiar properties of the object, which makes it difficult to transfer the same code to another object.

In the previous section we saw that there is an open question within neuroscience over whether the brain is practically medium dependent or independent, meaning that neuroscientists debate over whether models of the brain at a high level of abstraction (i.e. omission of details) can fully capture the brain's cognitive capacities. If the brain is practically medium dependent, no such description will be available. [19] The task of this section is to show why this is the more likely outcome. Prac-MD indicates a high reliance of the functionality of a system on the specific characteristics of its material constitution. The argument will be that prac-MD is the most thrifty way to code information and for this reason is most likely to have evolved.

No-one can dispute that resource consumption was a strict constraint in the evolution of nervous systems.[20] Information coding is energetically expensive because use of a physical system as a coding system means forcing it out of its thermodynamic equilibrium states (Polyani, 1968). Neurons are expensive cells to build and maintain. At the same time, large nervous systems confer survival advantages. The evolution of brains has had to tread the path that maximises information transmission while minimising costs. The only way to do that is to exploit the inherent coding and signalling capacities of the materials at hand (living cells, bundles of biomolecules). As Sterling and Laughlin (2015) argue, most of the signalling in the brain involves metabolically inexpensive, local chemical reactions, whereas more costly action potentials are needed for temporally precise, long range signalling. Evolution has ensured the efficiency of information processing in the brain by working with, not against, an unrestricted number of the physical properties of the material that the brain is made of. This is not compatible with achieving practical medium independence. To recall, prac-MI comes about when there is a formal description which maps on to a model of a concrete system at a high enough level of abstraction such that it can feasibly be implemented in a fundamentally different medium and retain the functionality of the original concrete system. This is possible if the fine and even medium-grained material details of the concrete system are irrelevant to its functionality. But if those details are irrelevant, it means that they are not being used for informational tasks, which means that the full coding potential of the physical system is not being exploited. Evolution could not afford to be so wasteful.

---

[19] And this would not be at all atypical amongst the physical objects that are modelled by scientists. It is normally accepted that the abstractions of models render them incomplete in their descriptions of the behaviour of their targets. The brain is an outlier in that ambitions for the completeness of computational descriptions have been set so high as to fully encompass the cognitive capacities of the brain.

[20] This is argued at length by Sterling and Laughlin (2015).

This argument draws inspiration from Conrad's (1989) comparison of the informational capacities of brains and programmable digital computers. Some of digital hallmarks that he notes, such as serial processing, are not present with neuromorphic architectures. Importantly, though, he points out that a digital computer is a "[p]hysical realization of a formal system" which means that the "equations of physics [are] irrelevant to computing", whereas the functionality of the brain is "[h]ighly dependent on material composition and physical interactions" (Table 2, p.209). Although this comparison is thirty-five years old, the accumulated body of results from neuroscience concerning the cognitive role of finely orchestrated molecular interactions at the synapses, dendrites and elsewhere only stands to confirm Conrad's observation.[21] Much of Conrad's argument concerns the restriction against greater utilisation of the inherent capacities of the material system due to the need for digital computers to be programmable. As indicated above, programmability demands that functionality cannot be heavily dependent on properties of hardware; if it did, re-programming would always involve tinkering with the hardware, as is the case with analogue machines. This leads Conrad to be skeptical about an equivalence between brains and digital machines:

> the brain could be replaced by a structurally programmable system with equivalent computational function only if the malleability of biological structure-function relations is completely irrelevant to adaptive behavior, an unreasonable expectation. (Conrad 1989:204)

The point is that in biology structures interact flexibly with function: functional plasticity is dependent on hardware plasticity, both within one lifespan and across generations. In order to better understand the brain we should move beyond notions inherited from digital computation, such as the separation of function from material structure, and come instead to conceive of information processing in the brain as being the result of a slow adaptive process in which the pre-existing physical-chemical properties of biological cells were leveraged to convey specialised codes.[22]

To achieve signalling on the cheap – limiting the costs of building materials and energy consumption – there are some basic principles that need to be employed, and these have the potential to better elucidate brain function and its evolution than the well-worn brain-computer analogy. Primarily, begin with what you have. This means co-opting existing molecular pathways of cells to do additional information processing tasks. This is an important point discussed by Cao (2022,17-19), about the close similarity between signalling cascades in the brain and metabolic processes. In general, a signalling system is cheap when it works with the existing dynamics of a physical system, rather than imposing a structure on it that conforms to an independently designed code or formal system – which is how a computer engineer might approach things. This is an observation made by Howard Pattee during his investigation of biological signalling, as summarized by Rączaszek-Leonardi (2012, 304-305):

> One might see formal symbol systems as a subtype of this more general class [of symbolic systems]—a subtype, in which the dynamics are artificially reduced to the point of being inconsequential…..Formal systems require *abstracting* from dynamics, an 'overdetermination' of the result, because they arose for very specific human purposes. This kind of formal precision is not required for cognition in the wild.

---

[21] See e.g. Larkum (2022). And see Milinkovic and Aru (2026) for a related argument on the difference between digital and biological computation.

[22] I make the point here with an example of coding via movements of a flock of sheep on a hillside: https://philosophyofbrains.com/2024/10/04/the-case-against-medium-independence.aspx

Efficiency is achieved when as many as possible of the physical properties of the system are put to service for signalling. The limitations on this exploitation of physical properties will be due to noise management – e.g. the need to regulate concomitant biological processes, prevent electrical cross-talk and unwanted chemical side reactions. But they will not be due to the demand for programmability by an external agent, or implementability in different hardware.

The upshot is that any system whose functions satisfy prac-MI conditions (1) and (2) will not be maximizing the information carrying capacities inherent to its particular physical medium. It cannot, because it needs to generalize across media. The greater the extent that the system is maximizing the information carrying capacities of its medium, the narrower the range of variant media will be, that can feasibly yield equivalent functionality, and this will certainly not be possible in fundamentally different materials. Hence, given any reasonable assumption about the resource constraints on evolved nervous systems, we can be confident that the brain is practically medium dependent.[23]

4. Taking Practical Medium Dependence Seriously

Here I discuss the implications of the brain's practical medium dependence, firstly for neuroscience, secondly for AI, and thirdly for the field of neuromorphic computing. The first point is that there is only a clear separation between cognitive and non-cognitive properties of the brain on the assumption that its functionality is practically medium independent – an assumption which can no longer be maintained. To see this, consider that in order for the functionality of a concrete computational or signalling system to be prac-MI, there must be a categorical demarcation between the set of its properties that map onto the formal description and ones that are irrelevant to it. This is what allows for duplication of a subset of the structural features of the system in another one with a fundamentally different material constitution. For example, a flashing light used for Morse Code is a medium independent signal in which the only feature relevant to the code is a difference in duration of the flashes corresponding to the dot and the dash. This same structural feature can be implemented as a difference in duration of sounds, or of lengths of marks in a notebook. All other properties of the physical media are code-irrelevant, though serving as background conditions which make the coding possible. In an electronic computer implementing a binary code, the relevant properties are the two specific voltage ranges that map onto 0 and 1. All else is irrelevant to the actual computation. In electronic computers, this demarcation is enforced by design. For example, within a range corresponding to the reasonable expectations of the user of a laptop, the voltages of the transistors will not be affected by moving it, fluctuations in its temperature, changes in atmospheric pressure, etc., at least not to the extent that the voltages fall outside of the ranges to be associated with 0 and 1.

Note that the enforcement of this demarcation adds additional costs in terms of resources taken to build and run a computer, for what it entails is not only that the majority of the properties of the physical system are irrelevant to the computation (therefore not exploited) but also that they need to be maintained steady or kept from affecting the code-relevant properties that do map on to the formal system. This means that resources have to be spent on keeping the irrelevant properties of the machine from affecting the relevant ones – an obvious example is the resource expenditure for building and running fans and other

---

[23] Note that none of this depends on or attempts to demonstrate the *optimality* of resource use in biological information processing. It could well be that these systems do not achieve a theoretical optimum. What I am trying to show is just that the prac-MI approach will be far inferior to the prac-MD one, with regards to efficiency. Biology cannot afford to be enormously wasteful, though some level of inefficiency might be tolerated.

cooling systems to prevent temperature from affecting the operations of electronic computers. The functionality of the brain, in contrast, is prac-MD. There has been no evolutionary pressure to enforce a separation between cognitively relevant and irrelevant properties, and given the costs needed to achieve such a separation, we can assume it does not obtain. As Godfrey-Smith (2016, 493) has put it, "[t]he line between the 'information processing' side of human brain activity and the metabolic side is porous." The implication for neuroscience is that the sub-disciplinary separation between cognitive and computational neuroscience, working at a high level of abstraction, and neurobiology dealing with cellular processes, should be taken as a matter of convenience, not reflecting a fundamental division between the cognitive and non-cognitive aspects of the brain.[24]

Stated bluntly, the second implication of the practical medium dependence of the brain is that human-like AGI running on digital hardware is not feasible. By human-like AGI I mean an AI that achieves the suite of integrated human capacities, including consciousness and understanding, by approximating the putative algorithms of the brain in a digital substrate. Because of practical medium dependence there will not be a formal description of the functionality of the brain which is rich enough to capture its known capacities but sparse enough that it can be mapped onto the physical processes of a system made from fundamentally different materials, and feasibly be implemented in that system. Of course, we already have digital simulations of some of the fine-grained molecular processes within the brain that are known to be cognitively relevant. The issue is that such processes are so inefficient to simulate in digital hardware that scaling up the simulations to anywhere near needed to encompass a neural circuit (let alone a brain area or whole brain) would be so resource hungry that it is unfeasible. This stands in opposition to the view that the currently dominant AI approach – deep learning in ever larger ANNs running on digital hardware – will eventually scale up to human-like AGI because it will effectively encompass and match the complexity of functional processes in the brain. This view is rejected on grounds that the abstraction to electrical signalling, which is the premise of ANNs, models only a tiny subset of the brain's cognitively relevant properties. Our conclusion does not rule out ANNs manifesting in the future, and already instantiating, an alien form of information processing that we may wish to call cognition or intelligence. But we should be under no illusion that they capture the essence of human intelligence, or biological intelligence more generally.

The third issue is whether human-like AGI can feasibly be implemented in neuromorphic hardware through a reduction in energy costs made possible via mimicry of some of the efficient strategies of biology. First it should be noted that the field of neuromorphic computing has tended to embrace the abstraction to electrical signalling just as much as the field of deep learning. However, within neuromorphic computing there has been some experimentation with more brain-like materials, such as (non-living) organic chemistry with membranes that have structures like ion channels (Najem et al 2018). [25] To date, neuromorphic engineering has not explored the living medium of neural tissue. But an implication of practical medium dependence is that biological cognition is inherently dependent on its substrate of living cells. This suggests that the only feasible path towards human-like AGI is through synthetic biology and artificial life (AL), though it is currently unrealistic to

---

[24] As also discussed in Chirimuuta (2022, 2024 chapter 4).

[25] Marković et al. (2020, 504) discuss the pros and cons of this strategy: "Another emerging lead in neuromorphic research is the use of fully organic materials, some of which may have the unique advantages of biocompatibility, and possibly low-cost fabrication through the use of printing technologies. Organic memristive devices are mechanically flexible, multifunctional and often low power. Their application to neuromorphic computing includes harnessing their electrochemistry to implement low-voltage artificial synapses and neurons. However, these devices tend to suffer from slow operation due to the low mobility of carriers in organic materials, high levels of variability, lack of reliability due to the instability of organic materials, and poor retention."

attempt the technical challenges that this would involve. That said, neuromorphic engineers can still benefit from paying closer attention to neurobiology in order to build more efficient computers, setting aside the question of AGI. Even if the full functionality of the brain can only feasibly be implemented in neural tissue, there may still be principles of efficient use of resources for information processing, discoverable in the brain, that are generalisable to other materials. This is the hope of Laydevant et al. (2024, 183) who write:

> the physics of our brains and bodies are very different from those of any substrate we are likely to use for present or near-future computers, including neuromorphic computers, so for neuroscience insight to impact how state-of-the-art AI is done, it cannot be too specific to the physical details of the implementation in animals….insights from neuroscience will be most impactful in AI if they are physics agnostic or are particularly amenable to implementation with semi-conductor electronics and photonics rather than biological matter.

Another approach, advocated by Harrison et al (2022), is to invest research effort into the materials science of life-like but not living *active matter*. This holds out the prospect of artificially intelligent soft robotics systems whose materials have some of the properties of living tissues that help to make biological cognition efficient and adaptive, without attempting full bio-mimesis.[26]

5. Summary and Conclusions

This paper was motivated by the need for a conceptual toolkit to assess the dependence of information processing in the brain on the material specifics of the brain, without begging the question in favour of radical independence, as we saw with the traditional concept of medium independence which can lead the debate into an empirically vacuous dead end. The notion of practical medium dependence/independence was introduced to this purpose. It was shown that practically medium independent computing does not optimize material and energetic resources and for that reason it is highly improbable that it evolved. The efficiency of practically medium dependent systems explains how it is that brains consume so little energy in comparison with digital AI, and also leads to the expectation that human-like AGI will not be feasible using digital hardware.

There is scope for more work on the specifics of the case that energetic limitations rule out human-like AGI in digital computers. For example, estimates could be made as to how computationally intensive models would need to be in order to simulate chemical as well as electrical information processing in a neural circuit. It would then be possible to determine if the feasibility constraint is closer to a limit due to energy available under current technological conditions, or to an impossibility grounded in the cosmic scale of resources that would be required. It would then be possible to know whether hypothetical analogue or quantum computers could potentially approach the problem because less demanding in energy than digital machines. In any case, I hope to have shown with this paper the philosophical interest of bringing practical questions of feasibility to bear on the seemingly abstract topic of medium independence.

---

[26] See also Kaspar et al. (2021).

Acknowledgments

I would like to thank audiences at The University of Glasgow, the ISPSM meeting in 2024, the Avenues of Intelligence workshop in Heidelberg 2024, the NeuroExp meeting in Berlin in 2025, and at the Columbia Center for Theoretical Neuroscience for their incisive discussions of this work. I also benefitted greatly from discussion with Federico Adolfi, Guillaume Dumas, Juanma Parrilla, and from the comments of two anonymous reviewers.

References

Aaronson, S. (2013). Why Philosophers Should Care About Computational Complexity. In *Computability: Turing, Gödel, Church, and Beyond*. MIT Press.

Anderson, N. G., & Piccinini, G. (2024). *The Physical Signature of Computation: A Robust Mapping Account*. Oxford University Press.

Arkhipov, A., Gouwens, N. W., Billeh, Y. N., Gratiy, S., Iyer, R., Wei, Z., Xu, Z., Abbasi-Asl, R., Berg, J., Buice, M., Cain, N., Costa, N. da, Vries, S. de, Denman, D., Durand, S., Feng, D., Jarsky, T., Lecoq, J., Lee, B., … Koch, C. (2018). Visual physiology of the layer 4 cortical circuit in silico. *PLoS Comput. Biol.*, *14*, 11.

Bowers, J. S., Malhotra, G., Dujmović, M., Montero, M. L., Tsvetkov, C., Biscione, V., Puebla, G., Adolfi, F., Hummelc, J. E., Rachel F. Heaton, B. D. E., Mitchell, J., & Blything, R. (2023). Deep problems with neural network models of human vision. *Behavioral and Brain Sciences*, *46*, e385.

Cao, R. (2022). Multiple realizability and the spirit of functionalism. *Synthese*, *200*.

Chalmers, D. J. (2014). Uploading: A Philosophical Analysis. In *Intelligence Unbound: The Future of Uploaded and Machine Minds*. John Wiley & Sons.

Chirimuuta, M. (2021). Your brain is like a computer: function, analogy, simplification. In *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*. Springer.

Chirimuuta, M. (2022). Comment on Neurocognitive Mechanisms by Gualtiero Piccinini. *Journal of Consciousness Studies*, *29*, 7–8.

Chirimuuta, M. (2024). *The Brain Abstracted: Simplification in the History and Philosophy of Neuroscience*. MIT Press.

Conrad, M. (1989). The brain-machine disanalogy. *BioSystems*, *22*.

Drayson, Z. (2025). Defending the medium-independence of computation. *Mind and Language*, *40*.

Godfrey-Smith, P. (2016). Mind, matter, and metabolism.*Journal of Philosophy*, *113*, 10.

Greshko, M. (2022). Nuclear Power for AI: How to reopen Three Mile Island safely. *Nature*, *634*.

Harrison, D., Rorot, W., & Laukaityte, U. (2022). Mind the matter: Active matter, soft robotics, and the making of bio-inspired artificial intelligence. *Frontiers in Neurorobotics*, *16*.

Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. MIT Press.

Haugeland, J. (1997). What is Mind Design? In *Mind Design II*. MIT Press.

Haugeland, J. (1998). Analog and Analog. In *Having Thought: Essays in the Metaphysics of Mind*. Harvard University Press.

Hooker, S. (2020). The Hardware Lottery. *arXiv*, *https://doi.org/10.48550/arXiv.2009.06489*.

Jamadar, S. D., Behler, A., Deery, H., & Breakspear, M. (2025). The metabolic costs of cognition. *Trends in Cognitive Sciences*, *https://doi.org/10.1016/j.tics.2024.11.010*.

Kaspar, C., Ravoo, B. J., Wiel, W. G. van der, Wegner, S. V., & Pernice, W. H. P. (2021). The rise of intelligent matter. *Nature*, *594*.

Kersten, L. (2024). An idealised account of mechanistic computation. *Synthese*, *203*.

Kulvicki, J. (2016). Recording and Representing, Analog and Digital. In *Giving a damn: Essays in dialogue with John Haugeland.* MIT Press.

Kuokkanen, J. (2022). No computation without implementation? A potential problem for the single hierarchy view of physical computation. *Synthese*, *200*.

Larkum, M. E. (2022). Are Dendrites Conceptually Useful? *Neuroscience*, *489*.

Laydevant, J., Wright, L. G., Wang, T., & McMahon, P. L. (2024). The hardware is the software. *Neuron*, *112*.

Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, *155*, 1.

Maley, C. J. (2021). The physicality of representation. *Synthese*, *199*.

Maley, C. J. (2023a). Analog Computation and Representation. *British Journal of Philosophy of Science*, 74, 3.

Maley, C. J. (2023b). Medium independence and the failure of the mechanistic account of computation. *Ergo*, *10*.

Maley, C. J. (2025). Declaring independence from medium independence. *Mind and Language*, 40.

Marković, D., Mizrahi, A., Querlioz, D., & Grollier, J. (2020). Physics for neuromorphic computing. *Nature Reviews Physics*, 2.

Najem, J. S., Taylor, G. J., Weiss, R. J., Hasan, M. S., Rose, G., Schuman, C. D., Belianinov, A., Collier, C. P., & Sarles, S. A. (2018). Memristive Ion Channel-Doped Biomembranes as Synaptic Mimics. *ACS Nano*, *12*.

Pareda, A. E. (2014). Electrical synapses and their functional interactions with chemical synapses. *Nature Reviews Neuroscience*, *15*.

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

Piccinini, G. (2020). *Neurocognitive Mechanisms: Explaining biological cognition*. Oxford University Press.

Piccinini, G. (2022). Neurocognitive Mechanisms: Some Clarifications. *Journal of Consciousness Studies*, 29, 7–8.

Pigliucci, M. (2014). Mind Uploading: A Philosophical Counter-Analysis. In *Intelligence Unbound: The Future of Uploaded and Machine Minds*. John Wiley & Sons.

Polger, T. W., & Shapiro, L. A. (2023). The Puzzling Resilience of Multiple Realization. *Minds and Machines*, 33.

Polyani, M. (1968). Life's Irreducible Structure. *Science*, *160*.

Rączaszek-Leonardi, J. (2012). Afterword. In *Laws, Language and Life: Howard Pattee's classic papers on the physics of symbols with contemporary commentary*. Springer.

Rooij, I. van, Guest, O., Adolfi, F., Haan, R. de, Kolokolova, A., & Rich, P. (2024). Reclaiming AI as a Theoretical Tool for Cognitive Science. *Computational Brain & Behavior*, *https://doi.org/10.1007/s42113-024-00217-5*.

Sala, A., Lizarraga, A., Caminiti, S. P., Calhoun, V. D., Eickhoff, S. B., Habeck, C., Jamadar, S. D., Perani, D., Pereira, J. B., Veronese, M., & Yakushev, I. (2023). Brain connectomics: time for a molecular imaging perspective? *Trends in Cognitive Sciences*, 27, 4.

Sarpeshkar, R. (1998). Analog versus Digital: Extrapolating from Electronics to Neurobiology. *Neural Computation*, *10*.

Schuman, C. D., Kulkarni, S. R., Parsa, M., Mitchell, J. P., Date, P., & Kay, B. (2022). Opportunities for neuromorphic computing algorithms and applications. *Nature Computational Science*, 2.

Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, D., Dean, M. E., Rose, G. S., & James S. Plank. (2017). A Survey of Neuromorphic Computing and Neural Networks in Hardware. *arXiv*, *https://doi.org/10.48550/arXiv.1705.06963*.

Simon, H. (1969). *The Sciences of the Artificial*. MIT Press.

Sprevak, M. (2018). Triviality arguments about computational implementation. In *Routledge Handbook of the Computational Mind*. Routledge.

Sterling, P., & Laughlin, S. (2015). *Principles of Neural Design*. MIT Press.

Thagard, P. (2022). Energy Requirements Undermine Substrate Independence and Mind-Body Functionalism. *Philosophy of Science*, *89*.

Williams, D. (2025). Two senses of medium independence. *Mind and Language*, *40*.