

The Case for Medium Dependence

Comment on *Neurocognitive Mechanisms* by Gualtiero Piccinini

By M.Chirimuuta (m.chirimuuta@ed.ac.uk)

Forthcoming in *Journal of Consciousness Studies*

Abstract

This commentary focusses on Piccinini's claim that neural signalling uses medium-independent vehicles, a claim which is crucial in the arguments given for the computational theory of cognition. I argue that Piccinini overestimates the role of spiking in the brain, and underestimates the significance of chemical signalling within and between neurons. There are important implications for how we should understand sui generis neural computation.

Neurocognitive Mechanisms (Piccinini 2020, henceforth NM) will set the agenda for debates over mind, brain, and modelling for a good while. There are many things to write about in response to such a rich and comprehensive work. I will zoom in on just one issue – the position staked out in the book, that neural signalling uses medium-independent vehicles. It is crucial for the arguments given for the computational theory of cognition. Here is the “Argument from the Functional Organization of the Nervous System”:

1. Neurocognitive processes are functional manipulations of medium-independent vehicles in accordance with a rule.
2. Functional manipulations of medium-independent vehicles in accordance with a rule are computations (Chapter 6).
3. Therefore, neurocognitive processes are computations. (NM:213)¹

In what follows I will mostly be taking issue with Premise 1, but I will also have some things to say about Premise 2 at the end.

¹ See also the “Argument from Semantic Information Processing” (NM:215) where Premise 2 holds that neurocognitive systems transduce information from different sensory organs into medium-independent vehicles.

Various reasons are given for the claim that neural signalling is medium-independent (NM:214). Firstly, that spiking is the dominant form of neural signalling. Secondly, that action potential arrival timing and firing rate are the properties of spike trains functionally relevant to cognition, noting that these properties are similar throughout nervous system. Thirdly, that these properties can be realized in physical media other than neural tissue. This leads to the fourth point, that the vehicles of neurocognitive processes are multiply realizable² and therefore medium independent.²

I will begin by making the case that Piccinini's position over-estimates the role of spiking. Spiking is without doubt the physiological indication of information processing carried out by the nervous system that has garnered the most attention in experimental neuroscience. But there is reason to think that this electrical form of signalling floats above a wide sea of chemical signalling, both intra- and inter-cellular. By chemical signalling I mean the processes involving molecules within the nervous system undergoing specific reactions and structural changes in response to the presence of other molecules, when those processes can reasonably be described as transmitting information. In their book *Principles of Neural Design*, neuroscientists Sterling and Laughlin illustrate the concept schematically (fig. 1). What's depicted here is that the Enzyme signals the presence of Substrate, A to the Receptor by breaking down A into two products, one of which induces a chemical change in the Receptor. For the Receptor, there is a reduction in uncertainty about the presence of A, and therefore Shannon information is transmitted.

² "The second kind of MR results when components of different types are organized in the same way to exhibit the same capacity. By components of different types, I mean components with different capacities or powers" (NM:56). Since the capacities or powers of the components are different, the overall system is not tied to one kind of physical medium, hence medium independence.

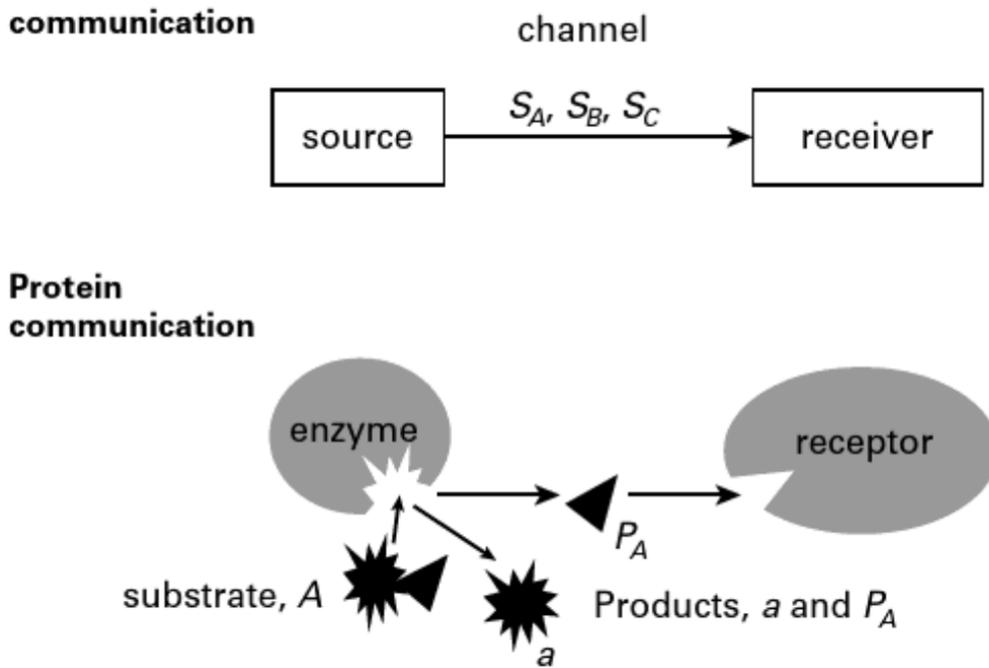


Figure 1. Protein communication modelled as a Shannon sender-receiver system. From Sterling and Laughlin (2015, Figure 5.1).

Sterling and Laughlin argue, furthermore, that considerations of the energy budget required for electrical (spiking) versus chemical signalling give in principle grounds for thinking that chemical signalling will be the dominant form of information processing in the nervous system. On their calculations, chemical signals consume many orders of magnitude less energy than electrical signals (2015: 35), and so on the plausible assumption that energy efficiency is one of the constraints governing evolution of the nervous systems, one should expect chemical signalling to be employed wherever possible.³ Conditions where this is not possible are when signals need to be transmitted over a longer distance than the scale of cells, and in a short enough time to be functional for control of the behaviour of a mobile animal.⁴

³ See Laughlin, de Ruyter van Steveninck, and Anderson (1998) on the energetic cost of neuronal spiking.

⁴ This is how Sterling and Laughlin (2015:105) summarise the view:

“as pulses transfer information over distance, they are mainly reporting results. The actual processing of information occurs mostly on a 1,000- fold finer spatial scale, the scale of molecules. There information is processed by chemical reactions: molecules diffuse, bind, exchange energy, change conformation, and so on. The key actors at this level are single protein molecules (~6 nm). They are targets for diverse inputs, such as small ‘messenger’ molecules that, upon binding to a receiver protein, reduce its uncertainty about a source.”

Lest the impression is given that the downward estimation of the role of inter-cellular electrical signalling is an idiosyncrasy of Sterling and Laughlin, I should mention that this is a point that has come up recently in various discussions of learning and memory. The classic account of memory, the one presupposed in artificial neural networks, is that learning occurs when the synaptic connections between spiking neurons either gain or weaken in strength, thus altering the flow of information through the system. This view is challenged by findings of mechanisms for adaptive, experience dependent alteration of behaviour – I will not get hung up on the semantics of whether this is ‘memory’, properly speaking – in single celled organisms, such as *paramecia*, which of course have no neurons. This disrupts the view that memory and learning in animals are due only to modulation in spiking patterns, through synaptic plasticity, because these intracellular mechanisms for adaptation are all pervasive and ancient, still occurring in the brain, and therefore must play a role in animals that also have synaptic plasticity.

This is how Ginsburg and Jablonka (2019, 462) make the case:

Neurobiologists have realized that they need to consider not only synaptic memory but also the epigenetic memory embedded in the cell nucleus and the transmission of ‘memory molecules’ between neurons (and other cells); these memory molecules include small RNAs, which can alter both the nuclear epigenetic memory within neurons and the synaptic connections between neurons. Thus, there are additional biochemical memory systems in the nervous system, not just the well-studied synaptic one, and although these mechanisms interact, they can be semi-independent, with nonlinear interactions. If we want to understand learning and memory in neural organisms, we need to consider all these mechanisms.⁵

⁵ Cf. Gershman et al. (2021, 2), “The possibility of a cellular-level mechanism for storing acquired information with delayed behavioral consequences is exciting from an evolutionary perspective because it suggests that the mechanisms for memory storage in complex multicellular organisms may have been inherited from much simpler organisms, possibly even protozoa, that share the same intracellular molecular repertoire.”

And even when it comes to synaptic plasticity, it has become clear that the molecular processes involved are far more complex than is required under the view that up and down regulation of synaptic strength is sufficient for learning. An implication of the discovery of the surprisingly large number of proteins expressed at dendrites is, Grant (2018, 3) argues, that the synapse is itself a “highly sophisticated molecular computer”. Again, the synaptic proteins are found to be evolutionarily very ancient, predating the emergence of neurons.

This closes the case for the importance of chemical signalling. I will now argue that this form of information processing cannot be medium-independent. The basic point I am making here is that for the vehicles involved in chemical signalling – the molecules – it just does not make sense to think of them as medium-independent vehicles. The molecules are signalling in virtue of their specific material (i.e. chemical) properties, such as binding affinities, and their modulating nanoscale structures, which determine how they operate and interact with other molecules. The medium is the message.

To properly establish this position, we must examine Piccinini’s criteria for medium-independence. He writes,

Medium-independent vehicle means that such vehicles are defined solely in terms of degrees of freedom and their organization, without reference to their composition or other concrete physical properties of the medium that realizes them. NM:129

What this requires is a clear cut separation between the physical properties of the vehicle that are and are not relevant to the representational content. In essence, my point is that you will not find this separation in the vehicles of chemical signalling in the brain and elsewhere. Put differently, the structural properties of the molecules cannot be separated from their functional ones. I suspect that the box and arrow diagrams normally used to represent signalling cascades give the misleading impression that these networks are made of swappable functional components – that all that is needed to specify the signalling process is to know the input-output function of each box, leaving questions aside about the actual role

fillers for each box, and their structural properties. However, the actual networks are elaborate, highly regulated sets of chemical reactions. They work as they do because of the specific chemical properties of the molecules involved. The box and arrow diagram is an abstraction away from material details, and it is there for the convenience of the scientist, but is in its own way a misrepresentation.

I have spoken in favour of taking chemical signalling to be medium-dependent. In addition, I do not think we should take it for granted that spike signalling between neurons is medium-independent. In electrophysiological experiments, the arrival of action potentials is literally heard as a series of clicks on a loudspeaker. Listening in on spiking neurons is a bit like hearing a telegram arriving in Morse code. These man-made codes are obviously medium-independent. Perhaps the resemblance fosters the idea that there is some equivalent code for neurons, such that all the information transmitted down axons is coded in rate and incidence times of the action potentials. Electrophysiologists listen in on the spike trains, but the receiving of signals within the brain, by post-synaptic neurons, is of a totally different kind. Action potentials arriving at axon terminals cause vesicles to be released, which leads to diffusion of neurotransmitters across the synaptic cleft; some of these molecules will bind to receptors on the dendritic spines and affect the membrane potential of the downstream neuron. The lesson here is that the inter-neuron communication mediated by spikes is not just electrical, it is chemical as well. So how can we say that only spike rate and timing matter to this form of information processing? The identity of the neurotransmitters and the receptors clearly matter as well. The same spike train (in terms of its rate and timing properties) will have a different effect on the downstream neuron – i.e. it will convey a different content – depending on the kind of neurotransmitter released, and the nature of the dendrites and their of receptors. And it is not that these additional features simply determine whether or not the spike train is effective in altering the membrane potential of the downstream neuron. It is a lot more complicated than that, and so we cannot presume that the entire content of inter-cellular signalling is conveyed by temporal structure of spike trains, in this Morse code like way.

To summarise, Piccinini’s argument for medium-independence of the vehicles of neurocognitive processes rests on the assumptions that spiking is the dominant form of signalling in the brain, and that rate and timing properties of are the only ones relevant to inter-neuronal signalling. Both of these assumptions are dubious. Still, Piccinini does acknowledge that there is more to signalling in the brain than spike trains – that neurotransmitters and hormones are other kinds of vehicles manipulated by neurons -- but he holds that “[a]nalogous considerations” for medium-independence would apply in those cases (NM:214). We should note here that neurotransmitters and hormones are molecules involved in inter-neuronal signalling, but the same considerations in favour of medium-dependence would apply there as with the intra-cellular “protein communication” discussed above. At this point in the book, Piccinini does not elaborate on what the analogous reasons are, but an indication is there in a passage from earlier on, used to exemplify the notion multiple-realizability₂ (MR2), which is a sufficient condition for medium-independence. He describes how circadian rhythms are found in an array of different species, and while the actual constituents of the system (the clock genes and clock proteins) differ, the organization is the same, involving oscillators and a transcriptional/translational feedback loop (NM:57-8). My thought here is that it would be highly speculative and implausible just to assert that MR2 occurs for all or even just for the majority of the signalling networks that scientists like Sterling and Laughlin are referring to. I am doubtful that many these networks (and there are very many of them) could have functional equivalents using different sets of molecules. In any case, the burden of proof is on Piccinini to show this, in order to defend the idea that chemical signalling is medium-independent.

I will now move on to discuss Premise 2 of the argument for the computational theory of cognition, the one that refers to the position defended earlier in the book, and by Piccinini (2015), that “[p]hysical computation is the functional manipulation of a medium-independent vehicle in accordance with a rule” (NM:129). This defines computation in a “generic sense”, which includes digital,

analogue, and sui generis neural computation. Whereas the vehicles for digital computers are digits, the vehicles for analogue computers are “physically implemented continuous variables”, and the vehicles for neural computation are (mostly) spike trains, these are all taken to be medium-independent.⁶ The thing I would like to point out here is that the demarcation between digital and analogue computation is matter of ongoing debate,⁷ and that on other proposals we find that medium-independence is a notion tied to digital computing in particular. This opens the door to ways of construing analogue and sui generis neural computation as operations over medium-dependent vehicles.

I have in mind Haugeland’s proposal that digital computers are *interpreted, automatic, formal systems*. He explains that,

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)

Every formal system is a digital system (1985:57). Digital codes are ones in which there is a “positive method” for reading or writing tokens: either it is completely successful or it fails. It is the black-and-white nature of digital coding which means that the content of a digital signal can be exhaustively determined and duplicated in another physical system, in any physical medium so long as it has the right number of degrees of freedom and so can preserve the structure of the code. Thus, for Haugeland, the requirement of medium-independence is tied specifically to digital computation, whereas he drops this condition for analogue computers – these are, “interpreted automatic systems that aren’t formal systems” (1985:259).

⁶ The “notion of medium independence can be used to characterize all the vehicles of computation, whether digital or not. Specifically, all physical computing systems—digital, analog, or what have you—can be implemented using physically different variables so long as the variables possess the right degrees of freedom organized in the right way and the implementing mechanisms can manipulate those degrees of freedom in the right way. In other words, computational vehicles are macroscopic variables defined in a medium-independent way. This, then, became my mechanistic account of computation in the generic sense: a physical computing system is a functional mechanism whose function is to manipulate medium-independent variables in accordance with a rule defined over the variables.” (NM:144)

⁷ See Katz (2016) and Maley (forthcoming)

The difference between formal and informal systems can be grasped intuitively by considering the difference between games like chess and games like football. With chess, any physical medium can instantiate the game so long as the right number of tokens and potential relationships between them can be distinguished. Likewise, chess boards and pieces can be made up of all kinds of weird stuff, in all sorts of sizes, or just played in a computer. With football, the nature of the physical medium in/on which it is played is integral to the game itself. It is not the same game if played on concrete instead of grass. An alteration to the dimensions of a grass pitch literally changes the game. A computer version of a football game could only be a mere simulation, not an actual game – whereas with chess, it is still real chess if you play it on a computer. The first part of this paper was making the point, basically, that most of the information processing happening in the nervous system is more like a game of football than of chess: what it is cannot be divorced from the actual stuff it is made of.

However, one way to push back on the claim that there is something importantly different about neural signalling, compared to signalling in a digital computer, is to invoke an equivalence between digital and analogue systems. Not that Piccinini or I are claiming that neural computation is analogue, it is still relevant to point out that analogue codes can be digitised – what amounts to the same information can be conveyed in either format – and so, the thought is, nothing too profound or interesting follows from my claim that neural signalling is medium dependent and therefore not digital (on Haugeland’s definition). If the information processing that goes on in the brain could still be digitised – i.e. duplicated in a digital machine – then most of the standard implications of the computational theory of mind still hold, e.g. regarding the multiple realisability of biological minds in non-living substrates.

It is here that another work by Haugeland (1981/1998), “Analog/Analog”, is highly relevant. He agrees that man-made analogue coding vehicles, such as film photographs, can be digitized because it is straightforward to specify what the

relevant physical dimensions of the code are, and give a digital approximation to them – i.e. by specifying a pixel size, and a range of grey scales or RGB values that will convey the information at a satisfactory resolution. He calls them “second order digital”. But other kinds of analogue vehicles are, “messier and touchier than pure analog—second-order messy”. The example he gives is the metabolic system of the rat, in which, “[t]he millions of delicate hormonal balances, catalytic reactions, surface effects, and immunological responses, all interdependent in a biochemical frenzy of staggering proportions” (p.86). The reason for saying that such a system is not second-order digital is that there is no obvious way to identify just the properties of the system that need to be included in the digitised version, against some others that are irrelevant, and to specify the resolution needed for the digital approximation of those relevant properties. Now if we adjust this example to consider a chemical signalling network in the brain (which is very much like a metabolic cascade),⁸ we see that this is just another way of making the point that Piccinini’s condition for medium-independence – that there be a clean separation between code-relevant and code-irrelevant properties of the vehicles – does not hold.

To conclude, I’d like to ask, *why not embrace medium-dependent neural computation?* The more we attend to signalling in the nervous system other than membrane voltage spiking, the harder it is to say that we are dealing with manipulation of medium-independent vehicles. We have seen that there is at least one other philosopher of mind and computation who makes medium-independence only a condition of digital computation, and Piccinini holds, anyway, that neural computation is not digital but *sui generis* (NM:chapter 13). And so, one option is to say that neural computation is functional manipulation of a medium-dependent vehicle. While theories of medium-dependent computation are less well developed than for medium-independent computation, this is an option well worth exploring

⁸ The fact that these two sorts of bio-chemical systems have this similarity raises the concern that if we focus on chemical signalling in the brain, we won’t have an answer to the question of how specifically neural systems give rise to cognition, since chemical signalling cascades are pervasive beyond the nervous system. I think that this possibility is not to be feared: in fact, the conclusion to be embraced is that there is more continuity between cognition and general features of living systems than is apparent on the standard computational theory of cognition.

and so I am left asking how this move would alter the whole architectonic of *Neurocognitive Mechanisms* – the many other components of this wonderful, wide ranging book.

REFERENCES

- Gershman, Samuel J, Petra EM Balbi, C Randy Gallistel, and Jeremy Gunawardena. 2021. "Reconsidering the evidence for learning in single cells." *eLife* 10. doi: <https://doi.org/10.7554/eLife.61907>.
- Ginsburg, Simona, and Eva Jablonka. 2019. *The Evolution of the Sensitive Soul*. Cambridge, MA: MIT Press.
- Grant, Seth G. N. 2018. "Synapse molecular complexity and the plasticity behaviour problem." *Brain and Neuroscience Advances* 2:1-7.
- Haugeland, John. 1981/1998. "Analog and Analog." In *Having Thought: Essays in the Metaphysics of Mind*, 75-88. Cambridge, MA: Harvard University Press.
- Haugeland, John. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
- Katz, Matthew. 2016. "Analog representations and their users." *Synthese* 193:851-871.
- Laughlin, Simon B., Rob R. de Ruyter van Steveninck, and John C. Anderson. 1998. "The metabolic cost of neural information." *Nature Neuroscience* 1 (1):36-41.
- Maley, Corey J. forthcoming. "Analog Computation and Representation." *British Journal of Philosophy of Science*.
- Piccinini, Gualtiero. 2015. *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- Piccinini, Gualtiero. 2020. *Neurocognitive Mechanisms: Explaining biological cognition*. Oxford: Oxford University Press.
- Sterling, P., and S. Laughlin. 2015. *Principles of Neural Design*. Cambridge, MA: MIT Press.