Preprint. Please cite the published version located here:
https://philosophyofbrains.com/2024/10/03/10-13140-rg-2-2-35310-63041.aspx

It takes two to make a view go right

Commentary on *The Physical Signature of Computation: A Robust Mapping Account* by Neal Anderson and Gualtiero Piccinini in The Brains Blog. September 2024

Danielle J. Williams
Washington University in St. Louis
danielle.williams@wustl.edu

The Physical Signatures of Computation is the most "robust" mapping view that's ever hit the market. It is impressive in its detail and the careful attention paid to its characterization of both the physical system and the formal computational description—a true service to the philosophical literature. The book promises a "unified account of artifact and biological computation," but here's where things take a turn: after handling artifacts, the Robust Mapping Account fades from view and the Mechanistic Account of Physical Computation (in its neurocognitive form) emerges as the way to capture biological computation. To have a unified account, then, requires two *distinct* approaches to understanding the nature of physical computation. This raises questions about whether the Robust Mapping Account can be extended to neural computation, but most importantly, questions about the compatibly of these two views is questionable given the history of the Mechanistic Account and its underpinning framework.

Importantly, both accounts (the Robust Mapping Account + The Mechanistic Account) have been sold to us as independent answers to the *same* questions. In other words, they are competitors on the market: each offering distinctly different answers to what has been called the "individuation question"

> **Individuation:** What is the difference between physical systems that compute and those that don't?

This question has been significantly bound up with what Chalmers (1994) calls the "implementation question."

> **Implementation**: When does a physical system implement a computational structure?

This feature of the literature—the blending of these questions—has been a source of deep frustration for me over the years. Not only are these questions bound up with one another, but the resulting picture ends up being that a view like the Mechanistic Account has the resources to answer the implementation question, when it doesn't seem to. The debate has carried on this way for quite some time. However, with this new book, the intermixing of these questions has finally taken a turn for the worse: two answers to the same questions have been given. While there are some differences between The Mechanistic Account proper and the Neurocognitive Mechanistic Account, the mechanistic picture that underpins both views remains the same, meaning the Neurocognitive Mechanistic Account is an *application of* The Mechanistic Account within a specific context, two views that cannot be cleanly pulled apart.

Because the Robust Mapping Account is framed as answering the implementation question (bound up with the individuation question) and the Mechanistic Account is framed as answering the individuation question (bound up with the implementation question), the authors might consider rethinking the questions that they intend to answer or else it is difficult to understand how these views can be brought together, what questions we are supposed to take them as answering, and whether mappings are cut out to handle neural computations at all.

The Mechanistic account does something quite radical when it came to physical computation. It proposes that we can understand the nature of physical computation by looking to mechanisms rather than providing a mapping view, or a semantic account. I describe this view as the "strong mechanistic account" (Williams, 2023). It is "strong" because it is *wholly* mechanistic and it this wholly mechanistic picture that makes it so radical.

One of the features that makes the Mechanistic Account wholly mechanistic is its reliance on the concept of medium independence understood as a mechanism sketch. Medium independence is the idea that computational processes "can be defined independently of the physical media that implements them… That is, computational descriptions of concrete physical systems are sufficiently abstract as to be medium independent" (Piccinini, 2015, p. 122). This way of thinking about abstract computational descriptions takes them to be abstract in the sense of abstract-as-omission: "I will call a description more or less abstract depending on how many details it omits in order to express that a system possesses a certain property" (Piccinini, 2015, p. 9). What is very important to notice here (and what makes it radical within the physical computation literature) is that an abstract description—a medium-independent description—is not a *formal* description—it is not an abstract-as-abstracta description—instead, it is a stripped-down description of one and the same thing: the behaving system in all its complexity. Not much changes, as we will see, when this account is developed further to capture neural computation.

Now that we've established that the Mechanistic Account targets physical mechanisms exclusively by giving an abstract-as-omission account of what it is to be a computational mechanism, I want to be clear as to why it does not answer the implementation question. Something important to recognize is that medium independence does not allow us to "abstract away" to the formal computational description. Put differently, you cannot go from an abstract-as-omission description to an abstract-as-abstracta description. Put differently yet again: There is no stepwise ladder from the mechanism to the formal computation. This should be of no surprise. The very nature of the implementation question is that it asks for a theory that draws the relation between these two types of descriptions—if there were such a ladder, a mapping view would be unnecessary.

Unlike the Mechanistic Account, the Robust Mapping Account does give us a bonafide answer to the implementation question. In Chapter 2, the authors describe—in great detail—how we should understand physical descriptions. "Physical descriptions are descriptions of the configuration and dynamics of physical systems, defined as the carriers of physical properties… physical descriptions are articulated within fundamental physical theories" (p. 44). Later in the chapter (Section 2.4) the authors describe the formal definition of computing systems. They say, "computational descriptions are articulated within well-defined computational formalism… [a]n example of a computational formalism is that of finite state automata…" (p. 54). FSAs are a paradigmatic example of an abstract-as-abstracta formal computational description. These two descriptions—the physical system and the formal computation—each make up one side of the implementation relation. Most of the book involves describing how this relation should obtain. It makes good on its promise and I think it does a very nice job.

In Chapter 9 the authors asks whether neurocognitive systems perform physical computations. However, in this chapter they bring in the Mechanistic Account along with work that's done in *Neurocognitive Mechanisms* (Piccinini, 2020). After some discussion, the authors arrive at the conclusion that the narrow teleofunctions of neurocognitive systems are to yield outputs that stand in certain mathematical relations to inputs and internal states… and aspects of spike trains that make the most functional difference are rate of transmission and timing. They argue that rate transmission and timing are multiply realizable properties and therefore the vehicles of neural processes are multiply realizable in the right way and since the vehicles are multiple realizable vehicles, they are medium-independent. This leads to the claim that "biological systems whose processes are medium-independent and manipulate vehicles in accordance with a rule are biological computing systems" (Anderson & Piccinini, 2024, p. 250). Let's unpack this.

On this account, multiply realizability is a relation between higher-level and lower-level aspects where higher-level properties are an aspect of their lower-level realizers. In his 2020 book Piccinini argues "that MR occurs when the same type of higher-level property is an *aspect* of different types of lower-level properties that constitute different mechanisms" (Piccinini, 2020, p. 39). Contra Fodor (1965; 1968) and Putnam (1960), multiple realization, in this context, is understood in terms of abstract-as-omission descriptions. This is because an "aspect" for Piccinini is "just a *part* of a property." Thus, the relation between properties is construed in terms of "parthood," i.e., part-whole relations (Piccinini, 2020, p. 26). To connect this to what is said above, when the authors state that spike rate transmission and timing are "multiply realizable properties," they mean this in terms of abstract-as-omission which is why, when the authors say that these features are multiply realizable, they can also say that *they are medium-independent*.

What does all this mean? Well, it means that when it comes to neural computation, the authors do not seem to "integrate the framework introduced in earlier chapters with recent work in this area to push the debate forward" (Anderson & Piccinini, 2024, p. 233). Instead, the authors transition from answering the implementation question (that has also claimed to answer the individuation question) to giving an account of what it is to be a (neural) computational mechanism in terms of an abstract-as-omission description (that has also been claimed to answer the implementation question). This means that we get two full accounts in the book: one on mappings and one on mechanisms where the former addresses computational artifacts and the latter addresses biological computation.

REFERENCES:

Chalmers, D. (1994). On Implementing a Computation. *Minds and Machines*.
Fodor, J. (1965). Explanation in Psychology in *Philosophy in America*. Routledge.
Fodor, J. (1968). *Psychological Explanation*. Random House.
Piccinini, G. (2020). *Neurocognitive Mechanisms*. Oxford University Press.
Piccinni, G. & Maley, C. (2021). Physical Computation in *The Stanford Encyclopedia of Philosophy*.
Piccinini, G. (2015). *Physical Computation, A Mechanistic Account*. Oxford University Press.
Putnam, H. (1960). Minds and Machines in *Dimensions of Mind*. New York University Press.
Williams, D. (*forthcoming*). Two senses of medium independence.

Williams, D. (2023). *Implementation and interpretation: A unified account of physical computation,* ProQuest Dissertations & Theses Global. (2865991732).