Forthcoming in *Mind & Language*. Please cite published version.

SYMPOSIUM

Two senses of medium independence

Danielle J. Williams

Washington University in St. Louis

Correspondence

[Danielle J. Williams]

Email: [danielle.williams@wustl.edu]

1 Brooking Drive

Department of Philosophy

St. Louis, MO 63130

Abstract

The term "medium independence" has different meanings. One sense maps onto "abstract-as-abstracta" descriptions while the other maps onto "abstract-as-omission" descriptions. Both senses have been deployed when it comes to understanding the nature of physical computation. However, because medium independence is a polysemic term, the sense being used should be clearly stated. If the sense is not clearly stated, then those who wish to engage in debates regarding medium independence and physical computation run the risk of conflating different but related issues which can have consequences when it comes to thinking about physical computation across contexts.

KEYWORDS

medium independence, multiple realization, physical computation, implementation

1. INTRODUCTION

The term "medium independence" has been used in two distinct, but related ways across the philosophical literature, especially in debates centering around the nature of physical computation.

One sense in which it is used maps onto the familiar concept of “multiple realizability” and involves the relationship between an abstract description and a physical system. This formulation understands “abstract” as abstracta where some formal (often computational) description is used to describe some physical process. This first sense is directly related to what has been called “the implementation question” which asks about the relation between a formal computational structure and a physical system (Chalmers, 1994; Williams, 2023). The second sense has its origins within the context of a mechanistic explanation and has been most recently used to underpin the Mechanistic Account of Physical Computation (Piccinini, 2015). This second use understands medium independence as providing an abstract characterization of a physical system where certain details about that physical system have been omitted. The difference between these two senses of medium independence is that the former understands medium independent descriptions in terms of “abstract-as-abstracta” while the latter in terms of “abstract-as-omission.” This distinction is important when it comes to discussions centering around the nature of physical computation.

In this paper, I distinguish between these two senses of medium independence by describing their origins within the philosophical literature. I then argue that medium independence in the first sense is directly related to the “implementation question” which asks about the relation between a formal computational description and a physical system (Section 2). Medium independence in the second sense is relevant to mechanistic explanation and, I argue, the “individuation question”. The individuation question asks for the difference between physical systems that perform computations and physical systems that do not (Section 3). After drawing this distinction, I provide some discussion regarding how one might answer different questions about medium independence and physical computation by showing how the two senses can be addressed and paired in different ways within various contexts (Section 4).

2. MEDIUM INDEPENDENCE AS “ABSTRACT-AS-ABSTRACTA”

The first sense of “medium independence” which understands abstract descriptions as “abstract-as-abstracta” comes from Haugeland (1985). In *Artificial intelligence: The very idea*, Haugeland gives a theory of what he calls, “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, p. 58).

Haugeland supports this idea by drawing a distinction between formal and non-formal games, such as chess or pegboard solitaire, where you can play either game using pieces made out of relatively any physical medium. For example, because chess is a formal game, you could play chess from opposing rooftop penthouses using thirty-two radio-controlled helicopters and sixty-four local rooftops. While an outlandish example, the point is that there are some games that can be executed using pieces made from any physical medium regardless of their physical properties or dimensions. This is contrasted with non-formal games such as football or billiards where the games are sensitive to size, shape, dimensions, texture, and so forth—where the physical medium very much matters (Haugeland, 1985, p. 58). For example, if you change features of the football field such as its length or width, it will affect the game as the dimensionality of the field and its parts matter for the game itself.

Haugeland, though, is ultimately interested in the nature of formal systems within the context of artificial intelligence and computational psychology with the following idea in mind: "Brain cells and electronic circuits are manifestly different 'media;' but, maybe, at some appropriate level of abstraction, they can be media for equivalent systems…a computer mind might be as much a real (genuine) mind as helicopter (or computer) chess is genuine chess, only in a different medium" (Haugeland, 1985, p. 63). Formal systems, then, are like formal games in that the tokens are manipulated according to rules which can be described independently of any physical medium. On this account different systems may have "higher-level" sameness because they can be given the same formal description. Consider one of Haugeland's examples comparing a computer and a brain. According to Haugeland they have higher-level sameness in that at the formal level they can be described using the same formal, computational description while literally being made out of different physical materials (brain matter and silicone, for example).

Higher-level sameness is captured by a formal description where the formal description is abstract in the sense of abstracta. Abstract, in this sense, can be contrasted with the "concrete" where an abstract-as-abstracta description (a formal description) can be given irrespective of the concrete particulars.[1] Formal descriptions in this sense have a place within philosophy of mind, especially within the context of functionalism. Machine-state functionalism, specifically, asserts that we can describe mental states formally, in this same sense, without specifying the physical realizers of those mental states (Putnam 1960, 1967). In this case, the formal description is an abstract-as-abstracta description given in terms of computation.[2] To restate this idea in Haugeland terms, any system with a mind has the same "higher-level sameness" and because of this, we can describe minds using formal descriptions, while accepting that different minds are made of different kinds of mediums—they are medium independent.

Within the context of functionalism this idea maps onto the familiar concept of "multiple realizability." According to Putnam (1967) mental kinds (characterized computationally) are multiply realizable in that they can realized by different physical mediums. For example, one and the same formal description of a mind given in terms of a Turing machine can be realized by different physical types such as the human brain, an artificial system, or some other minded organism we have yet to encounter.

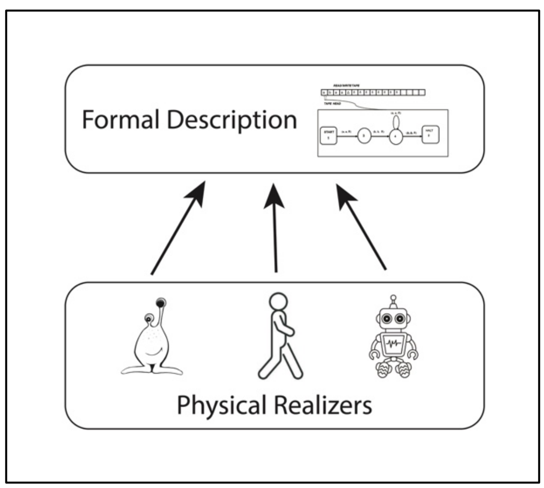

FIGURE 1 Multiple realization. The upward arrows in the image indicate that the description can be realized by different types of systems made of different types of physical mediums.

[1] But the question of how to draw the distinction between abstract and concrete objects is an open one: It is not clear how one should characterize these two categories nor is there a definite list of items that fall under one or the other category (Falguera *et al*., 2022).

[2] Originally as a finite-state-machine (Putnam, 1960) and then later as a Turing machine (Putnam, 1967).

Put into Haugeland terms once more, this is the idea that the multiply realizable computational descriptions are medium independent of the physical realizers. Thus, because Haugeland is interested in the medium independence of formal systems and computational processes are formal in this same sense, there is a direct connection between Haugeland's use of medium independence and the concept of multiple realizability. Call this sense of medium independence, "MR medium independence."

> MR medium independence: Formal (abstracta) descriptions can be specified independently of any physical medium.

Medium independence, then, in this sense involves a formal, abstract-as-abstracta description of some physical system given independently of their physical details.[3]

2.1 MR medium independence and the implementation question

There are different questions that we can ask about the nature of physical computation (or computation in physical systems). One question, "the implementation question", asks when a physical system implements a particular formal computational description (understood as a computational structure) (Chalmers, 1994; Williams, 2023).[4] An answer to this question is directly related to MR medium independence because it asks about the relation between the formal computational description and a particular physical system—the specific physical realizers of that description. In Figure 2, the downward arrow captures the nature of the relation.

[3] This is not to say that physical details will not constrain the types of formal descriptions that can be given, so there always is some connection between the formal and the physical such that the physical system must be capable of performing the formal description ascribed to it.

[4] Some examples include Turing machines, Pascal programs, connectionist networks, and finite state automata.

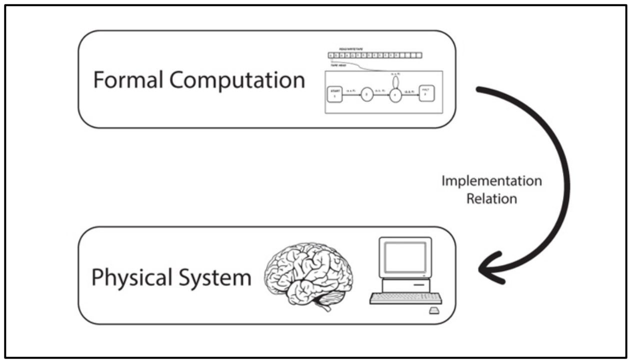

FIGURE 2. The implementation relation. The top box represents the formal computational description while the bottom represents the physical system. The downward arrow captures the implementation relation that obtains between the formal computation and the physical system.

In general, the implementation question has been answered by providing some variation of a mapping between features of the formal description and features of the physical system.[5] The mapping is meant to *connect* the computational description with the physical system such that we can say that the physical system *implements* that particular formal computational description.

An answer to the implementation question is proposed to have direct import into the computational cognitive sciences where computational descriptions of brain processes are commonplace (Chalmers, 1994). It is within this context that the implementation question becomes relevant to the functionalist framework.[6] The downward arrow in Figure 2 specifies the relation between the (multiply realizable) formal computational description and a particular physical system (the realizers). So here, we can see that multiple realization, functionalism, and the implementation question come together such that, in this context, the formal computational (abstract-as-abstracta) description of the mind is related to the physical cognitive system via the implementation relation.[7]

3. MEDIUM INDEPENDENCE AS "ABSTRACT-AS-OMISSION"

---

[5] Various theorists spell out this relation in different way (see Piccinini & Maley, 2021 for an overview).

[6] In fact, Putnam provides the very first mapping view. We can track one of the first theories of implementation to a body of work where Putnam describes functionalism and multiple realizability (Putnam, 1988).

[7] It is sometimes argued that non-mapping views, or *semantic* accounts answer the implementation question. While there is reason to think that they do not (Williams, 2023, chapter 3), I set these views aside for the purposes of drawing the distinction between different senses of medium independence and their relation to the computation literature.

Describing computation as "medium independent" has been popularized in the physical computation literature by the Mechanistic Account of Physical Computation, championed by Gualtiero Piccinini (2015). The sense in which medium independence is used within the mechanistic account comes directly from Garson (2003) (Piccinini, 2015, p. 122). Garson describes Adrian's 1928 theory of how we should understand information transmission in the electric activity of neurons. This conception is based on the generality of nervous transmission: Basic principles of information transmission in neurons obtain among several types of sensory receptors, biological taxa, and among motor neurons. The ubiquity of the structure of the nervous impulse implies the medium independence of the sequence: The temporal structure of the sequence can be instantiated across a wide range of physical mechanisms.

This core idea is used to develop the mechanistic account of physical computation. According to the mechanistic account, a process, such as performing a computation, is medium independent if it can be realized by different kinds of physical mediums irrespective of the higher-level properties that they may interact with. Careful attention should be paid to what is meant by "higher-level" properties because it differs from how Haugeland understands it. Higher-level properties within the mechanistic account are not formal descriptions (abstracta). Instead, they are stripped-down descriptions of one and the same thing: the behaving system in all its complexity.

Within this framework mechanisms are understood as being composed of entities and activities. Entities are things that engage in various activities and activities are the producers of change (Craver, 2001; Machamar *et al*., 2000). The organization of the entities and activities determines the way in which the mechanism produces the phenomenon. Functions, in this context, are the roles played by entities and activities in mechanisms. Within this picture, mechanistic explanations can be given at some level within a nested hierarchy where the lower-level entities, properties, and activities are components of the mechanism that produce higher-level phenomena. This picture provides an understanding of how mechanistic explanations can be given at varying levels of abstraction where, depending on the level in the hierarchy, more or less details about the physical system will be included (or omitted). To see this, consider that scientists do not always provide a complete mechanistic explanation. Instead, they sometimes provide what is called a mechanism "sketch" (Craver 2006; 2007). This places mechanistic explanations on a continuum depending on how much detail is provided by the description. One end of the continuum is the mechanism sketch: A description of a mechanism where bottom-out entities and activities cannot (yet) be supplied or where the explanation contains gaps in the stages of the mechanistic process. These gaps are present because there are missing pieces that scientists do not yet know how to fill in. Because there are missing pieces in the explanation, sketches are meant to indicate the further work that needs to be done in order to understand the mechanism in its entirety. A mechanism sketch that relies on abstraction—stripping away details—is what underpins Piccinini's concept of medium independence.[8]

With this idea in mind, we can understand medium independence in this framework as taking a kind of "bottom-up" approach to thinking about physical computation by starting with the physical system (a mechanism) characterized at some level abstraction where specific details about the physical medium—or structural details—are omitted. This approach can be contrasted with the above approach which starts with a formal computational description (abstracta) and then later relates that description to the physical system. This places the focus on features of the physical system specifically where the omission of the structural details involves a description of a mechanism that is medium independent in the abstract-as-omission sense:

> Medium independence: Some structural details are omitted from the mechanistic explanation (a mechanism sketch is provided).

The difference between medium independence and MR medium independence is that the former specifies features of a mechanism with some details omitted, while the latter specifies a formal computational description irrespective of the physical medium. MR medium independence is relevant to the question of how a formal computational description relates to a physical system, while medium independence is relevant only to the implementing mechanism.

3.1 Medium independence and the mechanistic account

The mechanistic account of physical computation is an account of what it is to be a mechanism that has the teleological function of computing. This views answers what is sometimes called "the

[8] I have characterized abstraction as the stripping away or omission of details. This simple description is to draw the contrast between formal descriptions and mechanism sketches, it is not to make a claim about the nature of abstraction as this is an ongoing debate. For example, see Levy (2021).

individuation question." This individuation question asks about the difference between systems that perform computations and those that do not.[9] According to the mechanistic account, what makes a physical system a computing system—the difference maker—is that it manipulates medium independent vehicles according to a rule (Piccinini, 2015, p. 10). According to this view we need not consider all of the specific properties of the physical system, only those properties that are relevant to the computation, according to the rules that define the computation.[10] For this view, a physical system can be described more or less abstractly, where "abstractly" is a description of a physical system that omits certain details—a mechanism sketch.

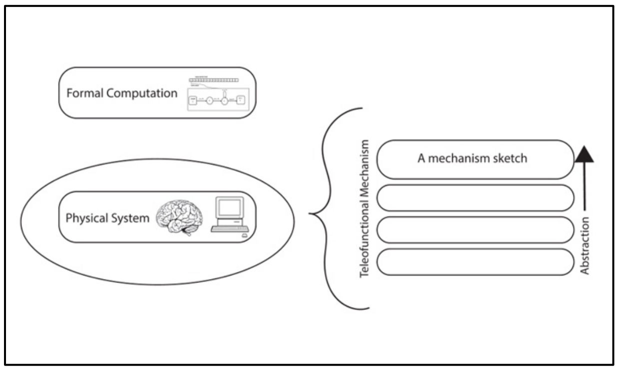

FIGURE 3 The mechanistic account of physical computation. The left-hand side of the image lacks a downward arrow demonstrating that the implementation relation is not being captured by the theory. The right hand shows how what is being captured are details about the physical system in terms of abstraction.

Importantly, though, the way medium independence is conceptualized within the mechanistic account makes up a critical feature of the conceptual apparatus that underpins the view; it is what marks it out from other theories that address the nature of physical computation. In particular, it contributes to the result that the mechanistic account should not be adopted as an answer to the

[9] It is sometimes claimed that theories of computational implementation are answering the individuation question (Piccinini, 2015; Shagrir, 2022). Here I keep these questions distinct and leave open whether theories of computiatonal implementation should properly be understood as also answering the individuation question. This is fine for the present purpose because whether implementation theories do or do not answer the individuation question has no bearing on whether the mechanistic account answers the implementation question as it is presented here.

[10] Rules are understood as an input-output mapping.

implementation question (because does not draw the relation between a formal computational, abstract-as-abstracta, description and a physical system) (Williams, 2023, chapter 7; 2024).[11]

4. THINKING ABOUT MEDIUM INDEPENDENCE

In sections 2 and 3, I characterized two senses of medium independence and their relation to different questions regarding the nature of physical computation. In one sense, medium independence is understood as abstract-as-abstracta where a formal description can be given of some physical system irrespective of its physical media (MR medium independence). Medium independent descriptions of this kind are directly related to the implementation question which asks for the relation between formal descriptions and a particular physical system. The second sense of medium independence is understood as abstract-as-omission where a mechanistic description of some physical system is medium independent when specific details about that system are omitted, specifically, details about the physical medium. In this section I will provide three ways of thinking about medium independence that touch on the relationship between the two senses I describe in the previous sections to illuminate the interesting interplay between the two senses when it comes to thinking about medium independence and computation across contexts, such as within cognitive science.

4.1 Rejecting medium independence while accepting MR medium independence

One question you might ask is whether every type of computational process is medium independent in the way that the mechanistic account proposes. For example, Maley (2021) argues that computation *in general* cannot be medium independent because analog computation is not

[11] The conclusion that the mechanistic account does not answer the implementation question follows Chalmers (1994) specification of the nature of the question. This point is argued for in detail by Williams (2023), but across the literature, all accounts that address the nature of physical computation in some way are typically described as answering the implementation question. Thus, the implementation question has become a kind of "catch-all" within the debate. However, drawing out the distinction between the different senses of medium independence in the literature does help to illuminate why the implementation question should *not* serve as a catch-all question when it comes to addressing the nature of physical computation. So, while I distinguish these questions here, it has not been standard practice in the literature to do so.

medium independent. Maley argues that representational properties of analog computations are part of the physical medium of those computations—that there is a physicality to the representations in an analog computer. This argument provides a counterexample to the idea that we can describe *all* computational processes as medium independent in the way that the mechanistic account proposes. To make this point, Maley draws on the example of a differential analyzer. Within the differential analyzer physical movements, such as rotation and displacement, are used to execute the computational task. This interconnected series of components computes solutions to mathematical operations, including differential equations. The parts of this mechanism are what do the representing and whether they are capable of representing certain variables in an equation has to do with their physicality—features such as their size, shape, and location. Maley's work targets medium independence in that the representational properties in analog computing systems are part of the physical system; meaning, we cannot omit them from the description in the way that medium independence requires within the mechanistic account.

Within the context of cognitive science, Maley argues that analog computation is a leading contender for understanding neural computation (Maley, 2022; 2024). Because he holds the view that computation is necessarily a representational process, neural computation requires that properties of representations are manipulated by mechanisms that are only sensitive to the properties that do the representing (Maley, 2022, p. 4). This requirement makes neural computation—analog neural computation—medium *de*pendent because neural representations are sensitive to the physical medium. This brings the mechanistic explanation "down" to a level that falls "lower" in the nested mechanistic hierarchy—one that includes details about the physical media that does the representing and that is no longer medium independent.

Maley argues against the medium independence of analog computation which challenges the mechanistic account directly when it comes to whether it can provide an account of physical computation that captures all types of computational systems. However, even if we accept this conclusion, it does not therefore mean that we have to also reject MR medium independence and the implementation relation. You can think that the there are some instances where specific computational processes won't be medium independent. However, you can still think that computational processes can be described MI medium independently. Put differently, it is possible for there to be a type of computational process that is not medium independent (analog computation) while also thinking that there are some types of physical systems that implement formal computational descriptions (e.g., digital computational systems). This shows that you can *reject* medium independence while *accepting* MR medium independence and a need for an answer to the implementation question.

4.2 Rejecting medium independence and MR medium independence

When it comes to neural computation, Piccinini argues that we can describe neural computations in terms of their spike rate and timing without saying anything about additional features of the mechanism (Piccinini & Bahar, 2012, p. 462; Piccinini, 2020). This is a claim that the appropriate level of abstraction for an explanation of neural computation is at the level of spike rate and timing. Piccinini arrives at this level of abstraction because he thinks that those features are medium independent features of neural processes and that they are the most relevant to neural computation. However, Chirimuuta (2022) has recently argued against this level of abstraction when it comes to understanding neural processes. In this critique, Chirimuuta argues that Piccinini overestimates the role of spiking in the brain and underestimates the significance of chemical signaling within and between neurons. This challenge rejects the idea that we can give a medium independent description of (computational) neural processes.

Ultimately, though, Chirimuuta argues in subsequent work that the brain should not be taken literally as a computing system (Chirimuuta, 2024). So, not only are neural processes not medium independent, the idea that the brain performs computations in general is also rejected. This second move is what, in turn, automatically rejects MR medium independence when it comes to computational descriptions of the brain and the resulting applicability of the implementation question. If the brain does not perform computations, then there is no reason to relate formal computational descriptions to neural processes.[12] Thus, it is possible, within a particular context, to reject both medium independence and MR medium independence when it comes to specific issues about physical computation.

4.3 Accepting medium independence and MR medium independence

Another option is to argue that computation is MR-medium independent while also denying that all there is to a theory that addresses the nature of physical computation is an account that describes

[12] As Chirimuuta (2024) argues, this does not entail that computational descriptions of the brain are of no use.

a medium-independent computational mechanism. This argument is an insufficiency claim about the mechanistic account of physical computation being taken on its own to provide an account that address *all* of the relevant aspects of physical computation. Put differently, you might argue that there cannot be a mechanism-only theory of physical computation (Williams, 2022). Such an argument says that to address computation in physical systems, part of what needs to happen is that there is also a specification of the conditions under which that physical system implements a formal computational description. So, not only should we understand what it means to be a computational mechanism, but we also need to understand how it is that the mechanism implements a formal computational description. This view accepts both medium independence and MR medium independence in that even if computation is medium independent, we still need a theory that addresses the implementation question.

5 CONCLUSION

I have described two senses of medium-independence and how they figure into theories about the nature of physical computation. The first sense has to do with a formal description, that is, abstract-as-abstracta descriptions while the second has to do with mechanistic explanation, that is, abstract-as-omission. The former sense is directly relevant to the implementation question while the latter, the individuation question. These senses are important to keep distinct, especially when it comes to thinking about issues related to physical computation. In the final section I explored the interplay between these two senses and I showed that it is not so simple when it comes to rejecting or endorsing medium independence. Questions about how the concept is embedded within the theory matter for understanding different ways of thinking about physical computation. This lesson should help philosophers thinking about the nature of physical computation across contexts, including cognitive science, think more clearly about how medium independence relates to physical computation and how different theories address different senses of the term.

Acknowledgements

I would like to thank Tom Polger and Larry Shapiro for organizing the symposium on computation and medium independence at the Central APA in 2024 where I provided commentary on the debate about whether computation is essentially medium independent alongside Corey Maley and Zoe

Drayson who each provided compelling arguments for adopting different perspectives on the medium independence of computation. I would I also like to thank Carl Craver for helpful comments on an early draft of this paper.

REFERENCES

Adrian, E. D. (1928). *The basis of sensation*. WW Norton & Co.

Chalmers, D. (1994). On implementing a computation. *Minds and Machines*, 4(4), 391–402.

Chirimuuta, M. (2022). Comment on neurocognitive mechanisms by Gualtiero Piccinini. *Journal of Consciousness Studies*, 29(7-8):185-94.

Chirimuuta, M. (2024). *The brain abstracted: Simplification in the history and philosophy of neuroscience*. MIT Press.

Craver, C. (2001). Role functions, mechanisms, and hierarchy. *Philosophy of Science*, 68(1), 53-74.

Craver, C. (2006). When mechanistic models explain. *Synthese*, 153, 355-376.

Craver, C. (2007). *Explaining the brain*. Clarenden Press.

Falguera, J., Martínez-Vidal, C., & Rosen, G. (2022). Abstract objects, In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*.

Garson, J. (2003). The introduction of information into Neurobiology. *Philosophy of Science*, 70(5), 926–936.

Haugeland, J. (1985). *Artificial intelligence: The very idea*. MIT Press.

Levy, A. (2021). Idealization and abstraction: Refining the distinction. *Synthese*, 198(24), 5855-5872.

Maley, C. (2021). The physicality of representation. *Synthese*, 199(5-6), 14725-14750.

Maley, C. (2022). How (and why) to think that the brain is literally a computer. *Frontiers in Computer Science*, 4, 1-6.

Maley, C. (2024). Computation for cognitive science: Analog versus digital. *WIREs Cognitive Science,* 15(4), 1-9.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.

Piccinini, G., & Maley, C. (2021). Computation in physical systems, In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*.

Piccinini, G., & Bahar, S. (2012). Neural computation and the computational theory of cognition. *Cognitive Science*, 37(3), 453-488.

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

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

Putnam, H. (1960). Minds and machines in S. Hook (Ed.), *Dimensions of mind* (pp. 1-20). New York University Press.

Putnam, H. (1967). *Psychological predicates in art, mind, and religion* (pp. 37–48). University of Pittsburgh Press.

Putnam, H. (1988). *Representation and reality*. MIT Press.

Shagrir, O. (2022). *The nature of physical computation*. Oxford University Press.

Williams, D. (2022, May 12-13). There cannot be a mechanism-only theory of computation. Rotman Graduate Student Conference: Models and Idealization. London, Ontario, Canada. https://www.youtube.com/watch?v=FcJ_oXuwudA

Williams, D. (2023). *Implementation and interpretation*: *A unified account of physical computation*. UC Davis. ProQuest ID: WILLIAMS_ucdavis_0029D_22222. Merritt ID: ark:/13030/m57n14c2.