Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Definition of Memory for the Cognitive Sciences

Permalink

https://escholarship.org/uc/item/11r0t01t

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Ross, Brett A.
Favela, Luis H.

Publication Date

2019

A Definition of Memory for the Cognitive Sciences

Brett A. Ross (brettross96@knights.ucf.edu)
Department of Electrical Engineering and Computer Science, 4328 Scorpius Street
Orlando, FL 32816-2362 USA

Luis H. Favela (luis.favela@ucf.edu)
Department of Philosophy and Cognitive Sciences Program, 4111 Pictor Lane, Suite 220
Orlando, FL 32816-1352 USA

Abstract

We provide a definition of ‘memory’ that is broad enough to apply to both natural and artificial systems. Inspired by computation and information theory, we define memory as a process that preserves information through time while maintaining its usefulness as an object to be computed. We defend the extensiveness of our definition by explaining how it applies to both brains and modern computers. We then consider potential objections to our definition. Our primary goal is to provide a definition of ‘memory’ that is broadly applicable across various cognitive sciences subfields.

Keywords: memory; computation; representation; information

Introduction

Memory is a central topic within the cognitive sciences and its various contributing disciplines, such as computer science, neuroscience, and psychology. One likely reason for this is its centrality to various conceptions of cognition. Be it brains or modern computers, memory typically plays a central role. However, it is often unclear if ‘memory’ is used the same across contexts. What is apparent, however, is the efficacy of computational theory in the cognitive sciences. Given the successes computational theory has provided the study of cognition, and given that memory is central to computation, it follows that a computationally-inspired approach to memory can provide useful insights into the general nature of memory. As such, it is necessary for us to explicate the relevant features of information and computation before discussing our definition of memory.

We begin with Piccinini and Scarantino’s (2011) definition of ‘computation’ as the processing of objects according to rules. Next, we connect that definition to Gallistel and King’s (2010) interpretation of Shannon’s classic information theory—that is, the reduction in uncertainty regarding the properties of an object—in order to show how computation allows for useful decisions about the world to be made. We then discuss some of the properties necessary for effective information processing and generic computation that describe modern computers, which may also be usefully applied to descriptions of brains as well. We pay special attention to the topic of representation. The ability of the definition to allow for determination of the boundaries of a computational system’s memory are examined. Finally, we present and respond to some potential critiques of the definition.

The definition these claims and terms are applied to is as follows: *Memory is a process that carries information forward in time, preserved in a fashion that maintains its usefulness as an object to be computed for the system to which the memory is said to belong*.

Defining Computation and Information

Before presenting our definition of ‘memory,’ we must first establish definitions for ‘information’ and ‘computation.’ We begin with information because, as will be discussed below, computation does not necessarily need to involve information—though it can be more useful when it does. Gallistel and King relate Shannon’s definition of ‘information’ as originating from a source, undergoing a process that ‘encodes’ the information into a ‘signal,’ and traveling to a receiver that ‘decodes’ the signal to derive a ‘message’ from it (2010, p. 2). The amount of information contained by the signal is determined not only by the signal, but by the receiver as well. The following example will make these points more evident.

Suppose an unseen coin is flipped and you are told, as a hint, that it might be heads or tails. You most likely already knew that and are wondering if this is really a hint at all. This highlights two important criteria for evaluating a signal’s informational content: First, a signal must be selected from a possible set of signals. How much information has been transmitted regarding an object depends on how the range of possible object states has been affected. The hint you received does not affect the range of possible outcomes from the coin flip, and thus holds no information. Second, the relative probability of the possible states under consideration plays an important role in evaluating the quantity of information transmitted. A coin is not a truly twodimensional object. There is a small possibility that it has landed on its side. The hint you received actually has some informational content, it is just small because the eliminated state is unlikely. Note that Shannon's definition of information does not restrict the types of objects and states that it describes. It may be something as quantitative as numerical data. Likewise, it may be something difficult to quantify numerically, such as the emotions of another. The key point is that there is a spectrum of possible properties and that the signal reduces their domain.

Computation invokes many concepts similar to information. In fact, as Piccinini and Scarantino point out, computation and information processing are often mistakenly held to be synonyms (2011, p. 3). We utilize Piccinini and Scarantino's definition of computation in general: "We use 'generic computation' to designate the processing of vehicles according to rules that are sensitive to *certain vehicle properties* and, specifically, to differences between different portions of the vehicles" (2011, p. 10; italics added). In the case that these vehicles are signals containing information, information processing *is* a form of computation as just defined. However, not all computation involves the processing of information. Informational content is not an intrinsic property of an object. It is relative to an observer and depends on how much the message reduces the observer's uncertainty (Gallistel & King, 2010, p. 7). Consider a computation that outputs 'cuidado' if the input is 'el horno esta encendido' and provides no output if the input is 'el horno esta apagado.' To an English-only-speaking observer, this computation does not process information—the objects have no meaning. But to a Spanish-speaker, this cautions them that the oven has been switched on. The computation performed is the same, regardless of the observer. Even if Spanish is forgotten, and the computation's objects cease to be meaningful to *anyone*, it is still the same computation. Thus, computation does not necessarily process information.

This definition of computation is clearly quite broad. It is so broad, in fact, that some philosophers believe that such an understanding of computation implies that everything performs computation, that is, 'pancomputationalism' (Piccinini & Scarantino, 2011, p. 5; cf. Chalmers, 2011; Copeland, 1996). It may be true that one could pick just about any physical phenomenon and find an arbitrary function that it computes (e.g., a rock; Chalmers, 1996). For "computation" to be a useful concept regarding research on cognition—such as memory—in the cognitive sciences, its scope must be appropriately pared down.

Recall that, based on the above definition, a computation is only sensitive to *certain* properties of objects, not necessarily all of them. A function that determines whether or not a neuron fires may only be sensitive to the firing/pre-firing properties of other neurons. Any additional physical variables are irrelevant to the purposes of the computation at hand, namely, modeling the dynamics of single-neuron activity. Pancomputationalism draws attention to the worry that "computation" may be a meaningless concept in research if it does not refer to some finite range of properties (or messages) that determine the results of the computations carried out by some system. In other words, its properties must have informational content that are relevant to the system.

With these conceptions of information and computation at hand, we can present a way to understand how they are present in the brain. Various brain processes can be usefully understood as computational, for example, the brain's ability to draw conclusions (Gallistel & King, 2010, p. 59). Consider the recognition of an image containing text (Figure 1). The optic nerve transmits visual stimuli to the brain, but it does not interpret the text's meaning. This is the role of a different portion of the brain. In this process there are signals (i.e., visual stimuli) that come from a set of possible messages (i.e., one image is distinguishable from another), which can be understood as processed in accordance to a set of rules that are sensitive to the signal's properties (i.e., the shape of the image is that of a word, and the word has meaning independent from the image). Here we have all the characteristics of computation being used to process information.

An important point to address is that of representation. In our discussion of information, we spoke of it as being encoded. In other words, it is represented within a certain syntactic structure. This encoding is what allows for reliable interpretation of the signal's contents. Modern computers contain a type of software called a "driver." Each driver instructs the computer how to interface with a certain type of peripheral device, such as a mouse or external hard drive. Despite the fact that both of these devices can communicate via a universal serial bus (USB) connection, the computer must use a very different set of rules when interfacing with a mouse than with a hard drive. Similarly, a computational description of the brain must refer to some syntactic structure when describing how the brain processes its signals. However, not any syntax will suffice. Both brains and computers are faced with a tremendous variety of possible objects to represent. A much simpler device than the brain is the TI-84 calculator. The largest number it can represent is approximately 10^{100}. If it was forced to have a unique character for each value, the number of unique characters would exceed the number of atoms in the known universe. The calculator avoids this conundrum by constructing its representations from a small number of symbols (i.e., 0/1 for binary, 0-9 for its decimal display) in a way that is sensitive to their relative positions. Similarly, the English language is represented through the use of twenty-six visual symbols (i.e., the alphabet) and forty-four audible symbols (i.e., phonemes) in a syntax that is sensitive to their relative positions in space and time respectively. All of these methods of representation are "compact," that is, the resources required to construct a representation grow logarithmically as the range of possible messages increases (Gallistel & King, 2010, p. 76). If even the humble TI-84 requires a robust syntax capable of compact representation, it follows that any

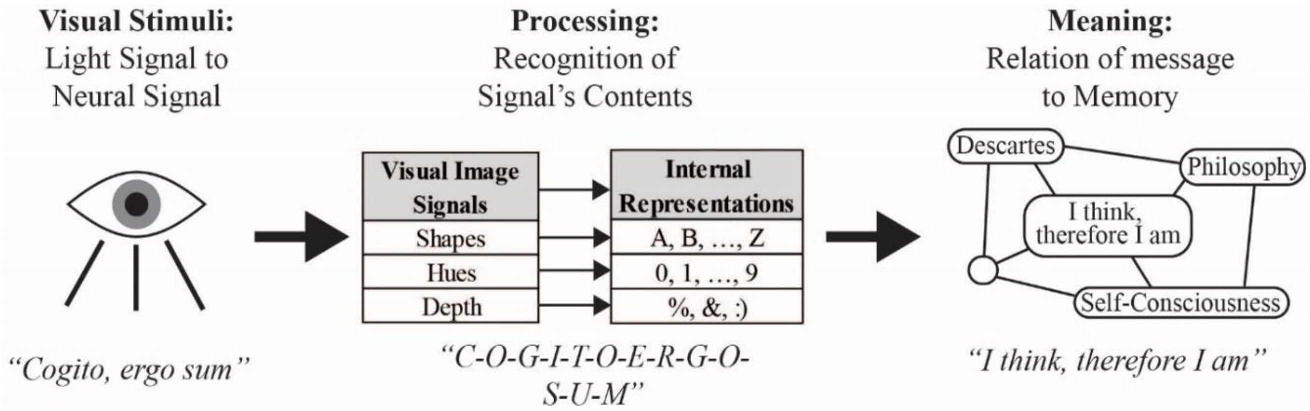

Figure 1: Reading text as the computational processing of information.

syntax that the brain may possess must also be capable of compactly representing the tremendous variety of messages it encounters (Gallistel & King, 2010, p. 82).

Like binary or decimal, the brain's syntax must be capable of constructing representations from a fairly small selection of basic elements. Otherwise, the incredible variety of sensory stimuli and the brain's practically unlimited creative capacity cannot feasibly be represented. These basic elements need not be numerical, and they need not be universal throughout the brain. Different portions of the brain perform different functions, and different syntaxes may be used. There will need to be commonalities between regions that facilitate their interaction, but there may be specialized computations (e.g., facial recognition, speech comprehension, etc.) that take place within regions. These computations might use a specialized syntax internally but use a shared syntax for computations that interact with other portions.

Having discussed computation, information, and representation, as well as what they mean in both the contexts of modern computers and the brain, we now move on to address how these ideas relate to memory and the role memory plays in the aforementioned contexts.

Memory in Computers and Brains

The property that distinguishes memory from other information-carrying signals in a brain or computer is its persistence through time. A fundamental requirement for a signal to be informative is that it is selected from a set of possible messages. There must be some way for this domain of possible messages to be established. When the layman enters an airliner cockpit, the array of dials and knobs are quite mysterious. To a trained pilot, each item denotes a meaningful piece of information. They know if an instrument's reading is alarming or typical. They know this because of their prior experience—information which was presented to them in the past and has persisted. In other words, their memory. Memory is the process that establishes the informational content of new signals.

Memory plays the same role in a modern computer. Without memory, a computer's only information regarding its past is that which is implicitly contained within its current state. As a result, whatever computation it performs must capture every relevant aspect of the computer's current state in order to determine the next state. Computers are often called upon to perform complex tasks that are combinations of a few basic functions (Gallistel & King, 2010, p. 109). Without a compact method for storing and preserving the results of past steps, any practical computation requires an absurdly large number of bits to define its state. Each instruction in a computer's program would have to shepherd hordes of bits. Suppose an instruction were as simple as providing a 1 (ON) or 0 (OFF) for each pixel in a display. To control every pixel of a typical 1080p resolution display, such an instruction would require *two million* bits. For comparison, modern central processing units (CPUs) typically use a humble sixty-four bits for their instructions. For a computational system to deal with these sorts of situations, a system is needed that has accessibility to the information contained in previous states as well as current—memory is needed (Gallistel & King, 2010, p. 131).

How does our definition of memory tackle this issue? In a modern computer, a CPU is the device that actually carries out most of the computations. In order to be able to perform computations quickly, it does not capture all of the information that it needs to perform all of its functionality within itself. If a CPU modifies an image, one cannot look at the CPU a few seconds later and determine how or what modification it performed. Instead, the CPU stores the image in a memory device to be retrieved later as needed. Later, when it needs to use the image in a computation, it calls upon the memory, which loads the image into the CPU. This transfer of bits constitutes a message to the CPU that informs it of the image's contents. The key point is that this is done without repeating the initial computation that resulted in the memory's message. Without the ability to call upon this *persistent information*, a modern computer would be as cumbersome as a pilot that has to be retrained for dial-reading every time.

Despite the similarity in the role of memory in brains and modern computers, there is a much greater degree of plasticity in the execution of the brain's memory processes. In certain environments, the way the brain remembers events can be highly vulnerable to suggestion (Loftus & Palmer, 1974, p. 588). This supports the idea that memory plays a role in establishing the informational content of new signals and highlights the dynamic nature of the brain's syntax. In Loftus and Palmer's experiment, use of the verb 'hit' versus 'smash' served to alter the subjects' syntax by priming them to think in certain terms. This in turn changed the message they obtained when they referenced their memory regarding the presence of broken glass at the scene of an accident. In this way, our definition addresses the inconsistency of the brain's memory—the syntax being used to interpret the stored information is constantly changing. These changes need not result in an insensible message. Rather, they result in a new interpretation. This differs from modern computers, where even slight changes in syntax can cause total malfunctions.

Explicating the neurobiological processes underlying memory in brains is not necessary for our project. The scope of our definition is readily understood in terms of Marr's three levels of description of information-processing systems (Marr, 1982/2010, p. 24). The first and most abstract level is the *computational theory*, which establishes the general feature of the system being investigated, such as vision, language, or memory. Next, is the *representation and algorithm* level, which describes the procedures for achieving said system feature. The final level is *hardware implementation*, which is concerned with the physical substrate forming the representations and carrying out the algorithms. In terms of explanatory strategies, these levels can be investigated individually. From this perspective, our definition is appropriately understood as working in the first two levels. Although we aim for our definition to be applicable to real systems, we leave work of explicating its physical implementation to others.

For our purposes, we merely note that if our definition of memory is appropriate for the cognitive sciences, then it can guide research that successfully identifies brain regions and processes that facilitate the kind of persistent information seen in modern computers as sketched above (cf. Srimal & Curtis, 2008). If our definition is incorrect, then there will be no empirical evidence of such persistent information. This follows from one consequence of our definition, namely, that in both the brain and in modern computers, memory serves the role of preserving information and establishing the possible set of messages from which new signals arise. In the next section, we explain how our definition of memory provides a way to delineate boundaries around the system in which memory occurs.

The Boundaries of Memory

For a memory to be computationally useful in the system it belongs to, a consistent syntax must be utilized during the encoding process—that is, preservation and representation—of signals. For example, the alphanumeric symbol '6' must always denote the quantity six, and not three or four. This stipulation helps establish who or what a certain memory belongs to. The boundaries of the physical system that consistently realizes a computational system's syntax then defines the boundaries of its memory.

As discussed earlier, a computation is sensitive to some properties of an object but not necessarily all. Specifically, it is reactive to particular forms of content, that is, information. When this computation handles information, the rules of its sensitivity must match up to the syntax in which the information is represented. This feature allows one to determine what contributes to a computational system's memory or not. In order to be memory, a process must not only carry some physical state forward in time, but the state it preserves must be preserved in accordance with the syntactic structure of that to which the memory is said to belong.

Modern computers possess a set of memory addresses, much like a set of street addresses in a neighborhood, that they have access to. Proper usage of these addresses is part of the syntactic structure of the memory process. Searching for an address outside of this range causes the memory process to malfunction. The signal the computer finds with such an address might be encoded using a different syntax, or there might not even be a physical signal present. Either way, if the sought signal is not represented in accordance with the syntax of the memory it is trying to find, it is not a part of the computer's memory. The signals found may inform the computer, but the information will not be accurate.

In many natural and artificial systems, it seems obvious where to draw the boundaries of—at least some of—their memory systems, for example, a human's hippocampus and a laptop's hard drive. In such cases, the syntax used by the computations are consistently applied only within the physical brain and hard drive. Accordingly, such memories are realized within an individual body or casing. With that said, our definition of memory is not a priori confined to brains and hard drives. As long as such features as information preservation and consistent syntax are maintained, the boundaries of memory systems are potentially quite broad. Though we aim here to apply our definition of memory to more traditional work in the cognitive sciences, we leave open the possibility of applying it to cases such as distributed cognition (e.g., shared remembering by couples; Harris et al., 2014) and cultural transmission (e.g., Rowlands, 1993).

The possibility of distributed or extended memory systems should not be controversial. A removable USB storage stick is external memory for any modern computer with a USB port and appropriate software. The memory is only available when the stick is plugged in, but a large number of different computational systems can all potentially access it. A written grocery list is external memory available to anyone who finds the list and can read the language it is written in (cf. Wagman & Chemero, 2014). The exact message the list presents will vary based on the individual's own internal memory—recall that the informational content of a signal is determined by the observer. Nevertheless, it is information carried forward in time. As long as the list's characters and words represent the intended message in a manner consistent with the reader's understanding of the language, the list can function as contributing to an external memory system.

We have attempted to show that our definition of memory is applicable to both narrow conceptions of cognition (e.g., isolated in brains), as well as more widespread notions (e.g., distributed cognition). Given that we discuss memory in terms of information processing, it is likely that the type of proponent of narrow conceptions who would readily accept our definition are those who think embodied, extended, and distributed cognition are still computational and representational in nature even if cognition is not isolated in brains (e.g., Barsalou, 2008; Hutchins, 1995; Wilson, 1994). On the other hand, it seems far less likely that anti-computational and anti-representationalists regarding cognition would accept our definition. We provide reasons why proponents of more "radical" conceptions of cognition could accept our definition by presenting experimental work involving affordances and memory.

Affordances are opportunities for behavior, and are based on the properties of the organism and environment (Gibson, 1979/1986). A doorway, for example, affords passing through for a human with narrow enough shoulders. Experimental work involving affordances stem from Gibson's ecological psychology (1979/1986). Contrary to representational approaches to perception, Gibson and his proponents argue that perception-action is not properly understood as centering on indirect representations. Visual perception, for example, is not a matter of an organism generating a mental image of the world, but instead is about an organism directly perceiving opportunities the world affords it.

Experimental work on affordances and memory have motivated conceptions of memory that do not appeal to computations or representations of the kind ecological psychologists and their proponents have resisted (e.g., Thomas & Riley, 2014; Vicente & Wang, 1998). Boschker, Bakker, and Michaels (2002), for example, conducted a set of experiments on the visual perception of climbing walls by experts and novices. When asked to recall information concerning the locations and orientations of holds on climbing walls, results suggested that experts can recall more information, clusters of information, and focus on functional aspects of walls (i.e., affordances); whereas novices did not recall clusters and focused on the structure of walls and not their functionals aspects. Boschker et al. argue that their findings show that differences in skill level correspond to differences in visual perception and memory. A central finding is that experts have memory that is better and of a more functional nature because they have more experience of perceived action possibilities than novices. In other words, their increased recall is tied to their increased perception of affordances. Note that this work does not appeal to computations or representations. Yet, our definition still applies: Experts have better task memory because the "information" relevant to action capabilities carries forward in time over the course of experience, and it does so in a manner that maintains its usefulness (i.e., affordance) to be "computed" (i.e., used) by the system (i.e., climber) for which the information belongs. The relationship between our definition and non-computational and anti-representational conceptions of cognition requires further fleshing out. However, we have attempted to demonstrate that the areas are not necessarily mutually exclusive. Having presented our definition of memory and discussed related issues, we now respond to several critiques.

Criticisms of a Computation-Based Definition

The appropriateness of utilizing our definition of memory in the cognitive sciences is contingent on the notion that it is explanatorily fruitful to describe the brain as performing computations. Computational approaches in the cognitive sciences are not without challenge. One source of opposition stems from forceful arguments claiming that phenomena investigated in the cognitive sciences are in no substantial way "computational," that is, "rule-governed manipulations of internal representations" (van Gelder, 1995). Therefore, our understanding of brains and cognition are set back by assuming they are like computers (Barrett, 2012). Another challenge centers on the claim that the prevalence of computationalism results from the prominent role of computers in modern society. Like other metaphors that were popular during their time, so too will the mind-as-computer metaphor pass (e.g., hydraulic pump, steam engine, etc.; Marshall, 1977). A third challenge is that many concepts underlying computational approaches have long and storied histories of imprecision. For example, many definitions of 'memory' now seem outdated in light of further technological advancement (Roediger, 1980). Addressing those challenges is far beyond the scope of the current work. Here, we respond to these criticisms in order to motivate the claim that complete rejection of a "computational" approach in the cognitive sciences is ultimately unwarranted.

First, unlike artifacts such as clay tablets or conveyor belts, computational theory is a set of formalized principles that are independent of any particular physical realization (Gallistel & King, 2010, p. 105). Computational theory becoming obsolete would be more akin to the obsolescence of calculus than that of the cellular phone. Computation is a field of mathematics, not a transient technology. While it is possible that computationally-based theories could be supplanted by non-computational ones (e.g., Chemero, 2011; Edelman, 1993; Kelso, 2009; van Gelder, 1995) for explaining all forms of cognition, such a shift would likely occur due to conceptual, methodological, and theoretical advances in the cognitive sciences, and not due to a technology's life-cycle.

Second, appealing to computational theory to investigate memory in both brains and computers does not necessitate that both compute digitally or numerically, or that their objects are both numeric or symbolic. Computational theory, in the form we appeal to, is consistent with identifying both modern computers and brains as "computational," even if the objects being computed differ. This is because what matters more than the realizers of particular processes is the syntax. Because modern computers manipulate digital objects (i.e., binary '1s' and '0s'), they are readily able to handle syntax involving computations of large digits, such as multiplication and division. Human brains, on the other hand, may not explicitly manipulate digital objects, which could account for the difference in speed of calculation. Specifically, brains may manipulate analog objects, which may not be as fast at processing syntax involving calculations of discrete numerical values. This would serve to explain why brains and modern computers have a different set of strengths and weaknesses and are better suited for performing different kinds of computation that are computations nonetheless.

A third reason to consider computation in some form is to appeal to the primary motivation for the cognitive revolution, namely, the need to posit "internal" states to more fully account for some kinds of cognition, action, and perception (Gardner, 1985). To be more precise, those cognitive capacities that occur without externally observable processes, for example, predicting and learning. Cognitive systems can make accurate predictions following very complex causal chains. An electrical engineer can look at a wiring diagram and tell what will open a certain contact without interacting with the real circuit. A complete explanation of this capability implies some internal process for simulating events and evaluating them according to a syntax. Cognitive systems can learn to perform behaviors without actually doing them. If a hobbyist reads an article on how to solder a wire before attempting for first time, they will certainly do better than if they had tried with no prior study. From these examples we do not further claim that cognitive systems are not embodied, that learning via action is likely necessary during developmental stages, or that physical practice improves abilities. Yet, such examples motivate the need to appeal to internal processes to fully explain some cognitive phenomena. In some cases, the most parsimonious explanation for these capabilities is the presence of internal representations and rules for consistent execution. All of this suggests that appealing to some form of computation to explain certain cognitive phenomena is well-motivated.

If we are correct that at least some cognitive capabilities (e.g., memory) are appropriately explained via internal processes of some sort, then the nature of how those processes represent must be accounted for as well. Although a tremendous deal of research and effort has gone into mapping and studying brain activity, there is yet to be evidence of a discernible syntax. This could be seen as evidence against computation in the brain. However, this may be a case of a lack of evidence not being evidence of absence. Gallistel and King explain that the more efficient and robust an encoding scheme is (i.e., representation), the less it resembles its message (2010, p. 4). The sheer variety of stimuli the brain is presented with suggests that its syntax would be extraordinarily complex, far more so than binary (e.g., neurons as on-off switches). Additionally, recall that it is not necessary for these representations to be discrete *or* numerical in nature. They might not even be expressible in terms of language. The brain has been produced by natural selection, not a highly-organized team of computer scientists. As such, there is no reason to believe that any criteria other than effectiveness for survival and reproduction has played a role in its development. There has been no force in natural selection pushing the brain's representations to be legible to outside observers. With all this in mind, it is no surprise that the brain's syntax remains a mystery.

The definition of memory posited in this paper proposes a broader definition of computation and representation than are typically applied to the brain. It also does not propose computation as an explanation of brain structure and function. Rather, it appeals to computation to describe memory processes. The aim of this is to enable a discussion that escapes some of the limitations traditionally associated with computationalism. This paves the way for the utilization of computation as a descriptive tool without rejecting other accounts of cognition (e.g., dynamical). While some systems are better explained by either computational or dynamical models (van Gelder, 1995), others benefit from the use of multiple explanatory strategies (Favela & Chemero, 2019). Depending on the goals at hand, one model may be preferable to another, and it is possible that neither can give an all-encompassing account of the system. Here, we are chiefly concerned with defining what memory is, and not the computations or dynamics that explain how it is realized in systems.

Conclusion

We have presented and defended a definition of memory. We began with Gallistel and King's (2010) formulation of Shannon information and highlighted the feature of observer dependence. We then presented Piccinini and Scarantino's (2011) broad conception of computation as the processing of objects according to rules sensitive to certain properties of those objects. Computational systems are distinguished from one another based on what properties they're sensitive to. If the objects being processed are signals with informational content, then the computational system processes information. Both the brain and modern computers can be described as such systems. Casting memory in terms of computation and information effectively describes memory as playing a role in establishing the meaning of new signals, that is, determining their informational content. The rules according to which these messages are interpreted are their syntax. In order to accommodate the wide variety of messages they represent, both the syntaxes used by brains and computers should be compact, that is, the resources required for representation should grow only with the logarithm of the number of possible signals. Having specific syntax for the purpose of carrying information forward in time allows for delineating boundaries around memory systems. We referred to ecological psychology's concept of affordances in order to illustrate that our definition is not necessarily incompatible with non-computational and anti-representational conceptions of cognition. We defended our use of computation as a tool for describing brain processes. Despite its challenges, computation's status as a set of formalized principles, as well as the ability of representations to serve as a succinct explanation of certain cognitive phenomena, make it well-suited for use as a descriptive tool. By limiting our use of computation to the *description* of memory, we remain nonpartisan as to the methods suitable for explaining its realization. As such, the following definition of memory is broadly applicable across the cognitive sciences: *Memory is a process that carries information forward in time, preserved in a fashion that maintains its usefulness as an object to be computed for the system to which the memory is said to belong*.

References

Barrett, L. (2012). Why behaviorism isn't Satanism. In J. Vonk & T. Shackelford (Eds.), *The Oxford handbook of comparative evolutionary psychology*. Oxford: Oxford University Press.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*, 617-645.

Boschker, M. S., Bakker, F. C., & Michaels, C. F. (2002). Memory for the functional characteristics of climbing walls: Perceiving affordances. *Journal of Motor Behavior*, *34*(1), 25-36.

Chalmers, D. J. (1996). Does a rock implement every finite-state automaton? *Synthese*, *108*, 309-333.

Chalmers, D. J. (2011). A computational foundation for the study of cognition. *Journal of Cognitive Science*, *12*, 323-357.

Chemero, A. (2011). *Radical embodied cognitive science*. Cambridge, MA: MIT press.

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

Edelman, G. M. (1993). Neural Darwinism: Selection and reentrant signaling in higher brain function. *Neuron*, *10*(2), 115-125.

Favela, L. H., & Chemero, A. (2019). *Explanatory pluralism: A case study*. Manuscript submitted for publication.

Gallistel, C. R., & King, A. P. (2010). *Memory and the computational brain: Why cognitive science will transform neuroscience*. Chichester, UK: Wiley-Blackwell.

Gardner, H. (1985). *The mind's new science: A history of the cognitive revolution*. New York, NY: Basic Books, Inc.

Gibson, J. J. (1979/1986). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Harris, C. B., Barnier, A. J., Sutton, J., & Keil, P. G. (2014). Couples as socially distributed cognitive systems: Remembering in everyday social and material contexts. *Memory Studies*, *7*, 285-297.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Kelso, J. A. S. (2009). Coordination dynamics. In R. A. Meyers (Ed.), *Encyclopedia of complexity and systems sciences* (pp. 1537-1564). Berlin: Springer-Verlag.

Loftus, E. F., &, Palmer, J.C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, *13*(5), 585-589.

Marr, D. (1982/2010). *Vision: A computational investigation into the human representation and processing of visual information.* Cambridge, MA: The MIT Press.

Marshall, J. C. (1977). Minds, machines and metaphors. *Social Studies of Science*, *7*, 475-488.

Piccinini, G., & Scarantino, A. (2011). Information processing, computation, and cognition. *Journal of Biological Physics*, *37*(1), 1-38.

Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition*, *8*(3), 231-246.

Rowlands, M. (1993). The role of memory in the transmission of culture. *World Archaeology*, *25*, 141-151.

Srimal, R., & Curtis, C. E. (2008). Persistent neural activity during the maintenance of spatial position in working memory. *NeuroImage*, *39*, 455-468.

Thomas, B. J., & Riley, M. A. (2014). Remembered affordances reflect the fundamentally action-relevant, context-specific nature of visual perception. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(6), 23612371.

van Gelder, T. (1995). What might cognition be if not computation? *Journal of Philosophy*, *91*, 345-381.

Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological Review*, *105*(1), 33-57.

Wagman, J. B., & Chemero, A. (2014). The end of the debate over extended cognition. In T. Solymosi & J. R. Shook (Eds.), *Neuroscience, neurophilosophy and pragmatism: Brains at work in the world*. New York, NY: Palgrave Macmillan.

Wilson, R. A. (1994). Wide computationalism. *Mind*, *103*, 351-372.