Embodied cognition as analog computation

Citation for published version:
Isaac, A 2018, 'Embodied cognition as analog computation', *Reti, saperi, linguaggi*.
<https://www.rivisteweb.it/doi/10.12832/92298>

Document Version:
Peer reviewed version

Published In:
Reti, saperi, linguaggi

Embodied Cognition as Analog Computation

Alistair M. C. Isaac

Abstract

Embodied, enactive cognitive science has traditionally rejected computationalism and its compatriot, representationalism. I argue that this rejection is too hasty, and places undo weight on intuitions about simple dynamical systems, such as the Watt Governor. I suggest instead that enactivists should consider cognition analogous to more complex, functional dynamical systems, such as analog computers, offering in particular a hydrodynamic computer, the MONIAC, as a new metaphor for embodied cognitive science. The implication of this approach is that adequate explanations of embodied cognition will require both the construction of models and the identification of representations.

§1. *Introduction*

Computationalism is the view that brains (and/or minds) are analogous to computers, and consequently the language and concepts of computer science provide a basic framework within which to construct psychological explanations. The fundamental implication of this analogy is that cognition comprises sequences of *transformations* over *representations*.[1] When this perspective first arose in the latter half of the 20th century, it contributed to the emergence of cognitive science as an integrated research program by facilitating communication across disciplinary boundaries—psychologists, neuroscientists, computer scientists, and philosophers could share ideas and debates by framing them in the language of computation. This communicative value in part explains why radical alternatives to the traditional perspective, such as connectionism or Bayesianism, are typically presented within the broad framework of computationalism.

Nevertheless, a growing nexus of related, high-level paradigms within cognitive science explicitly reject computationalism. A *locus classicus* for this rejection is van Gelder's (1995) "What might cognition be, if not computation?" Van Gelder argued that the dynamical systems approach to cognitive science suggests a new explanatory paradigm, requiring neither computation nor representation, yet able to provide deep understanding of the mechanisms of cognition. This challenge to computationalism has been echoed more recently in the work of the "multi-E"

---

[1] A recent trend in philosophy of computation argues that computations may be identified and their vehicles individuated through mechanistic analysis that makes no appeal to "representations" (Piccinini, 2008, 2015; Dewhurst, 2016; Schweizer, 2017). My point here is not about individuation, but constitution: the objects transformed during a computation are bearers of content, a presupposition of the mathematical theory of computation and all its results. This does not rule out the possibility of content-ignorant individuation of computations, but it does rule out the coherence, or at least relevance, of content-absent computational theories. Ironically, Piccinini himself seems to concede this point when granting that an "internal semantics" for computational vehicles follows from analysis of their functional role (2008, 135–6, 173–5; c.f. Isaac, *forthcoming*). The explanatory value of representation identification in analog computers is addressed explicitly in §3; those whose representation-phobia is still undamped by this discussion may seek solace in Miłkowski, 2015.

cognition movement, which emphasises the embodied, embedded, extended, and/or enactive nature of cognition;[2] Hutto and Myin (2013), for instance, explicitly reject both computationalism and representationalism. A somewhat more temperate viewpoint pays lip-service to computationalism, while rejecting the explanatory value of representations (Chemero, 2009; Villalobos and Dewhurst, 2017).

The purpose of this paper is to suggest that the shift to a dynamical, embodied, or enactive perspective does not obviate the computational metaphor. Rather, embodiment considerations imply that analog, rather than digital, computation is the appropriate model for explaining cognition. I argue that this point has been obscured by an undue focus in the dynamic cognition literature on the overly simplistic example of the Watt Governor (§2). I suggest instead the MONIAC, a hydraulic model of macroeconomic cash flow, as a new exemplar for computation in a dynamical system. Once we accept the MONIAC as a metaphor for embodied, enactive cognition, we see that the enactivist rejection of representationalism is misguided—*pace* Piccinini and Dewhurst, computation is constitutively the manipulation of representations; moreover, *pace* Chemero and Hutto, identification of representations in dynamical cognitive systems does substantive explanatory work (§3). Nevertheless, taking cognition to be analog computation does have novel implications, suggesting that representations are kinematic patterns rather than static symbols; dynamical explanations constitutively require simulations; and only robust, replicable neural processes are properly cognitive (§4).

§2. *Dynamics and anti-computationalism: The Watt Governor*

Enactive cognitive science emphasizes the role of dynamic interactions between brain, body, and environment in generating the complex behavior we identify with cognition (Varela, Thompson, and Rosch, 1991; Ward, Silverman, and Villalobos, 2017). By placing emphasis on these dynamic interactions, enactivism makes a radical break with the traditional, mid-20th century, research program in cognitive science, which focused on static representations, transformed in discrete time steps by syntactically specified rewrite rules, in an environment encapsulated from the world, such as the interior of a computer (or a brain). Enactivism replaces this static and discrete model with one on which the mathematics of cognition is that of dynamical systems theory, cognitive transformations are continuous, best described by differential equations rather than rewrite rules, and these equations may only be interpreted by assigning some variable values to quantities in the environment. Thus, questions about whether enactive cognition involves representation or computation turn directly on the appropriate interpretation of the mathematics of dynamical systems in the context of cognitive science.

Van Gelder (1995) defends the claim that dynamical systems models of cognition need not appeal to representation, *a fortiori* computation, through a detailed discussion of the Watt Governor. This ingenious device, invented by James Watt in 1788, used the centrifugal force of spinning weights to regulate the flow of energy in

---

[2] Multi-E cognitive science comprises a multitude of distinct research programs, not all of which are mutually compatible (Rowlands, 2009; Wheeler, 2017). Here, I coarsely group enactive and embodied tendencies on the grounds they are largely united in their rejection of computationalism, with the understanding that more nuanced distinctions are possible, and that not all within this tradition do explicitly reject the computational metaphor.

early steam engines. The problem is to ensure constant speed of a flywheel driven by the engine, and itself driving some further mechanism (wheels, loom, pile driver, etc.); the challenge is that the flow of steam from the engine fluctuates randomly, as does the consumption of energy, due to variance in task demands. For instance, if a steam train is to maintain constant speed while going over uneven terrain, the supply of energy must fluctuate up as hills are climbed and down as valleys are entered. Watt's solution was to establish a dynamical feedback loop between spinning weights driven by the flywheel and the flow of steam from the engine: as the weights' rate of spin goes up, the increasing centrifugal force pushes a lever closing a valve regulating the flow of steam; as their rate of spin decreases, the reverse occurs, and the drop in the weights' rate of spin releases the valve, allowing greater flow of steam. Consequently, the dynamic interaction between spinning weights and valve ensures flow of steam (and thus energy) maintain the desired equilibrium (Figure 1).

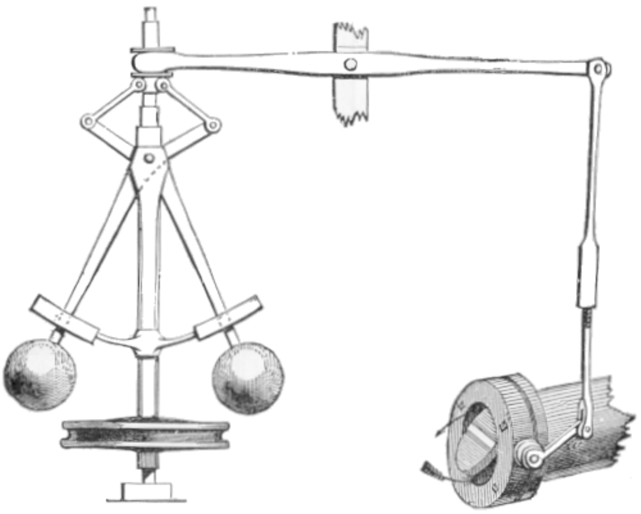

Figure 1: The Watt Governor.

In discussing this example, van Gelder contrasts it with an obvious (digital) computational solution to the same problem. A modern computer might regulate flow by (i) measuring and representing current quantity of steam; (ii) measuring and representing current speed of flywheel; (iii) representing desired speed of flywheel; (iv) performing a calculation to determine change in quantity of steam required to achieve desired speed from current; and only then (v) adjusting valve to change flow of steam. This analysis fits the computational paradigm: it stipulates a set of relevant representations (current speed, desired speed, etc.) and transformation over them, calculating the degree to which the valve must be adjusted by applying rules to the represented values.

Van Gelder takes representation and computation to be "mutually interdependent" properties (351), and thus, if the Watt Governor is not representational, it is inappropriate to describe it as computational. Yet the parts of the Watt Governor do not in any straightforward sense represent values of the sort found in the computational analysis of its task. The angle of the arms, for instance, might at first seem to represent engine speed; yet arm angle and engine speed are only directly correlated when the system is in stable equilibrium, while it is precisely when the system is out of equilibrium that the governor performs its function, dynamically adjusting steam flow toward the target (352–3).

A more important reason to reject analysis of the Watt Governor as computational or representational is that an effective description of its behavior—either the qualitative one above, or a more rigorous one provided by the differential equations it satisfies—does not appeal to representations or their transformation. Appeal to representation in explaining the Governor would be explanatorily idle: "representation is just the wrong sort of conceptual tool to apply" (353). This is the essential claim that later advocates of the dynamical perspective have emphasized: even if "a representational gloss is possible, once one has the dynamical explanation, the representational gloss does not predict anything about the system's behavior that could not be predicted by the dynamical explanation alone" (Chemero, 2000, 638).

Since van Gelder, the Watt Governor, and the question of whether or not it is representational, has been extensively discussed in the literature on embodied cognition (Chemero, 2009, 68–71; Shapiro, 2010, 119–25, 144–9; Hutto and Myin, 2013, 59–63). However, this discussion is relevant to cognitive science only if we accept that the Watt Governor serves as an adequate metaphor for embodied cognitive processes. Yet there is good reason to doubt this adequacy, namely the sheer simplicity of the Watt Governor: its internal behavior varies along a single functional dimension—height of weights while spinning—, as does its output behavior—opening or closing a valve. In contrast, complex cognitive behavior exhibits many dimensions of variance—uttering a sentence, ascribing a motivational state to a peer, coordinating with a stranger while carrying heavy furniture up stairs: all involve many degrees of freedom. Perhaps we should take a more complex dynamical model as our metaphor when considering the philosophical implications of embodied cognitive science.

§3. *A new metaphor: The MONIAC*

The field of analog computation studies complex, functional dynamical systems: the right sort of corral from which to wrangle a new metaphor for embodied cognition.[3] Most generally, an analog computer is a physical process functionally structured to constrain the relationship between continuously varying quantities. It is a computer in the sense that it realizes an effective procedure for transforming input values to output values that stand in a functionally specified relationship to its input and internal states. It is analog in that it deals in quantities that vary continuously, rather than discretely.[4]

Mathematically, analog computers typically solve systems of differential equations. Much as Turing (1936) broke digital computation down into a set of

---

[3] The claim that cognition is analog computation has been defended before by other routes, for instance description of neural dynamics (Rubel, 1985) or top-down functional analysis (Shagrir, 2010).

[4] It is important to stress that the quantities manipulated by an analog computer vary continuously from a *physical* perspective, as do weight, velocity, position, etc. This does not imply that they are continuous in the *mathematical* sense, or that they admit the arbitrary degree of specificity that can be found in a continuous mathematical object such as the real line. Confusion on this point has resulted in much nonsense about hyper-Turing computation (Isaac, Szymanik, Verbrugge, 2014, 790; Piccinini, 2015, 258–9, 271; c.f. §4.3).

fundamental operations in order to analyze its mathematical properties, Claude Shannon (1941) offered a model of analog computation in terms of a different set of fundamental operations: addition of curves, multiplication by a constant, and (crucially) integration.[5] Unlike digital computers, analog computers operate by physically *simulating* the desired functional relationships; for instance, in the Differential Analyzer, the particular computer inspiring Shannon, integration is performed by taking as input the motions of two variably rotating shafts *u* and *a*, representing functions of change over time, and constraining the motion of a third shaft to rotate at a rate varying as $\int_{t_0}^{t}(u + a)\,dt$ (338). The physical relationship between the rates of rotation of input and output shafts is analogous to, and thereby simulates, the mathematical relationship between change in curvature and circumscribed area of the target functions. For this reason, some presentations of analog computation emphasize not the continuous nature of the quantities involved, but rather the role of the system as an "analog" to the target of computation (e.g. Soroka, 1954).

There are many intuitive reasons to think analog computers might serve as an appropriate metaphor for embodied, enactive cognition. Perhaps the most fundamental is that analog computers satisfy the same mathematics as the dynamical systems that inspired van Gelder, Chemero, and the enactivists, namely systems of differential equations. Moreover, the theory of analog computation addresses the mechanical solution of differential equations with a very important constraint, namely that the solution procedure be *effective*, i.e. instantiable in a physical, real-world device that produces outcomes robust in the face of noise. This is important, because dynamical processes may only be targets of evolutionary selection insofar as they can repeatedly and reliable generate an adaptive behavior, i.e. perform robustly in precisely this sense. So, insofar as embodied cognition studies adaptive cognitive behavior, it should likewise restrict itself to positing mechanisms that exhibit this same robustness (c.f. §4.3). Finally, analog computation constitutively relies on continuous dynamics and naturally implements homeostatic feedback loops, exactly the kind of qualitative features that have inspired the multi-E cognition movement.

Lets look at an example of a specific purpose analog computer: the "Phillips Machine," more affectionately known as the MONIAC. The MONIAC was designed to model the cyclic flow of income in an open market economy by Bill Phillips, NZ electrical engineer turned LSE sociology student. He constructed the prototype in 1949, and its success earned him a lectureship at the London School of Economics; in the 1950s, Cambridge, Oxford, Manchester, Melbourne, Harvard, Chicago, the Ford Motor Company, and the Central Bank of Guatemala all purchased MONIACs. The MONIAC allows users to visualize the dynamic effects of policy on money supply by watching changes in the flow and pooling of colored water pumped through a sequence of pellucid tubes and tanks. (The water representing) income is pumped to the top of the machine, taxes funnelled off at the left, savings funnelled off at the right, then government expenditures and investment channelled back into the overall cascade of monetary flow, before expenses due to imports are funnelled off and, finally, income from exports flows back in to combine

---

[5] One of the complexities of the study of analog computers is the variety of possible formalizations, without the same clear convergence as was found across the methods of Turing, Church, and Post for digital computation. I gloss over these difficulties here.

with the remainder in the "transaction balances" tank at the bottom (Fig. 2). Interactions between these values are controlled through a system of valves and gates, which may be adjusted to represent, e.g., different rates of taxation or saving. Likewise, different theories of qualitative features of the economy (e.g. Keynesian vs. classical economics) may be modeled by holding tank levels fixed or allowing them to vary freely.[6]

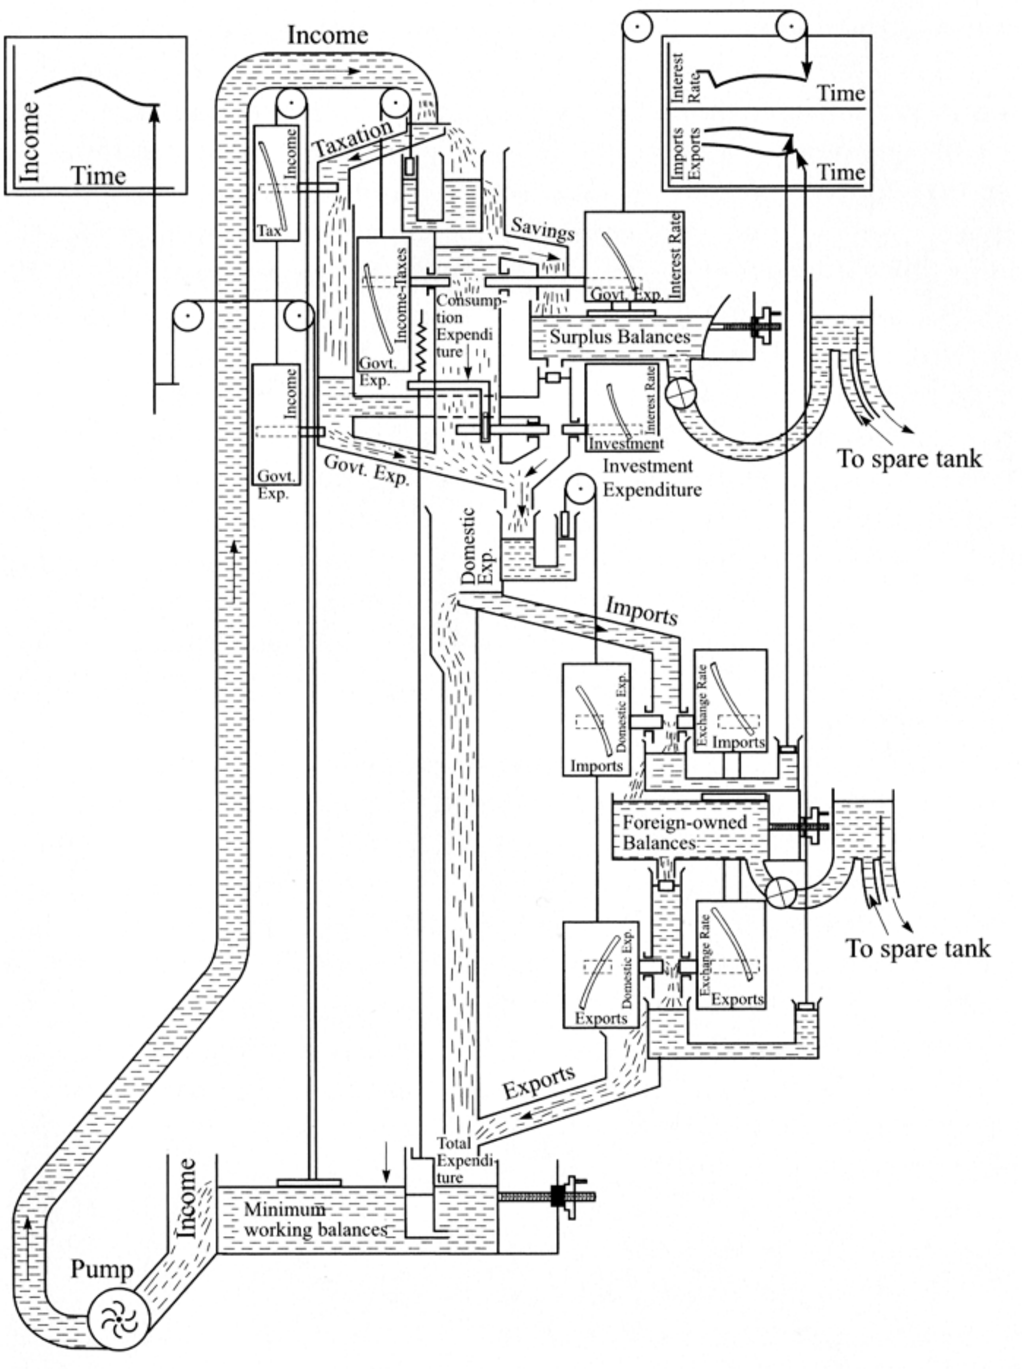

Figure 2: Schematic diagram of the MONIAC (reproduced from Barr, 2000, with permission of Cambridge University Press).

---

[6] This description draws heavily on Barr, 2000; see also Morgan, 2012, and Phillips, 1950.

The MONIAC was initially constructed as a teaching device; this explains Phillips' decision to use the visually enthralling (but mechanically temperamental) flow of colored water to represent income, rather than, for instance, change of voltage in an electrical circuit. Nevertheless, the operation of the MONIAC is precise enough to perform calculations (within a margin of ±4%), both of the effect of some policy intervention in establishing a new equilibrium between investment and expenditure, but also on the temporal trajectory of the change in monetary assets as the system transitions between equilibrium points. Floats in each of the tanks are connected to pens that plot the time change trajectory in the corresponding value for both visualization and quantitative analysis. Phillips' official (1950) description of the machine managed to combine a presentation of its mechanics and construction with a rich analogical discussion of its implications for monetary theory, arguably "the first application of dynamic control theory to macroeconomics" (Barr, 2000, 95).

Is the MONIAC an analog computer? Historians consider it a canonical instance of the genre (Bissell, 2007). It satisfies the general definition given above in that it constrains physical quantities (rate of flow of water, volume of water in holding tanks, positions of valves, etc.) in order to realize functional relationships between their values. Furthermore, it constrains these values by means of an effective procedure, bounded by an error margin of ±4%; within that bound, the performance of the machine is repeatable and robust. The localized mechanical interactions of the MONIAC realize basic operations analogous to those studied by Shannon: combining streams (e.g. when investment flows into the central column) implements *addition*; valve diameter implements proportional *multiplication* (as increasing diameter is equivalent to multiplying flow by a larger constant); and operations functionally equivalent to *integration* are implemented by those structural features that ensure variance in output with change in rate of flow, for instance the timeplots of change in investments or imports, or the rate of investment itself, which a clever spring-operated counterbalance ensures varies with the rate of change in income.

So, the MONIAC is an analog computer, but is it the right sort of device to motivate intuitions about embodied cognition? Above, I pointed out some general reasons to suspect enactive, embodied models of cognition are analogous to analog computers; one reason to consider the MONIAC in particular as a suggestive metaphor is that it was specifically designed to aid visualizability. Thus, while a general purpose analog computer, such as the Differential Analyzer, may more nearly approach the complexity of a human brain than the MONIAC, it is nowhere near as easy to visualize qualitatively. This easy visualizability ensures the MONIAC may serve as an effective prod to intuitions about analog computation. Another reason to embrace the MONIAC as a model for embodied cognition is that it shares key features with the Watt Governor: both Governor and MONIAC are equilibrium systems, maintaining their equilibrium points through homeostatic feedback loops. Dynamic feedback loops are suggestive for enactive cognition since they de-emphasize the boundary between cognitive system and world, highlighting instead the continuously fluctuating response by the agent to dynamic perturbations from the environment. Likewise, in the MONIAC, one may easily interpret any of the side channels to and from a holding tank as the contribution of an external "environment" to the homeostatic mechanism of the rest of the device.[7]

While the MONIAC implements the same kind of homeostatic feedback loops we see in the Watt Governor, it maintains a much more precarious equilibrium point, affected by many more variables. With this added complexity, come very different intuitions about the role of computation and representation in effective maintenance of dynamic equilibria. In contrast to the Watt Governor, representation appears to play a constitutive role in explanations of the MONIAC; in particular, assignment of representational role to tanks, valves, and channels is critical for explaining their structural differences. For instance, both the top and bottom tanks on the right have the same gross structural features, e.g. their levels may be held constant by linking them to an external storage tank. However, one must interpret the top tank as representing savings, and the bottom one as representing foreign holdings, in order to understand or explain the differential significance for overall dynamics of this feature—in the former case, holding levels stable is equivalent to fixing the interest rate, in the latter case, it fixes exchange rate. Moreover, recognizing these dynamics as instantiating a computation suggests a principled boundary between the functionally significant and the irrelevant aspects of performance; namely, only those features of the dynamics robust to fluctuations under 4% of overall flow rate participate in computation.

This last point allows us to see why mere identification of system dynamics (*pace* the anti-representationalists) can never constitute a complete explanation of a computation-performing analog mechanism. In a non-intentional dynamical system, such as global climate, no aspects of the system have privileged status—there is no difference in the type of explanation required for relatively stable features, such as mean rainfall in Chile's Atacama Desert, versus manifestly chaotic ones, such as hurricane formation in the Gulf of Mexico. It is this egalitarian quality which ensures that the differential equations which govern the system constitute a complete explanation of its behavior. In contrast, in a dynamical system that computes, some values are privileged over others: those values the system strives to maintain in equilibrium play a different functional role than those it does not; fluctuations in the system above the threshold of noise are functionally significant, while those below it are not. Correctly identifying these functional differences requires assigning a representational role to the values at stake, since only then can we assign correct functional role to the transformation or maintenance of these values. At the very least, this implies that some parts of the system will be *about* other parts of the system, serving to represent, and perform transformations over, those internal states.[8]

As a final illustration of these claims, consider the curved wall of the "surplus balances" tank at the top right of the MONIAC. The curvature of this wall ensures a particular functional relationship obtains between the volume and height of water in the tank, which in turn affects the rates of flow both in and out of the tank. If the

---

[7] Perhaps the most natural choice is the flow of import and export resources from the domestic portion of the model to the global economy. Phillips himself suggested that two MONIACs might be linked through these channels to study the effects of import and export policy between two nations (1950, 305). In enactive terms, the contrast is analogous to that between the kind of feedback loop that maintains equilibrium with a stochastically stable environment, and that established when coordinating with other embodied agents, such as when playing jumprope or football.

[8] This is the "internal semantics" of Piccinini, 2015.

anti-representational account of dynamical systems explanation is correct, an adequate explanation of this part of the system would simply consist in the differential equation governing this relationship. However, while such an equation would tell us what the curvature of the wall *does*—as might be adequate if this were a non-intentional system rather than a computer—, it would not tell us *why* the wall is curved, nor why the wall of the lower tank (in other respects, seemingly functionally very similar to the top tank) is not curved—the kind of understanding we demand of the MONIAC construed as a computational system. This difference can be explained by observing that the curve represents the functional relationship posited in Keynesian economics between the preference for liquidity, i.e. money available for immediate transactions, and the interest rate. It ensures the height of water in the tank is inversely proportional to, and thus serves as a measure of, interest rate. Likewise, the height of water in the lower tank measures the exchange rate; yet there is no analogous relationship between flow of money in and out of a domestic economy and anything corresponding to a notion of liquidity preference at the international scale. Only by going beyond mere description of temporal dynamics to an identification of represented values and represented functional relationships, are we able to adequately explain these features of the MONIAC.

§4. *Some morals for embodied cognition*

The previous discussion emphasized the reactionary consequences of taking the MONIAC as a metaphor for embodied, enactive cognition: *pace* the "radical" view, embodied, enactive cognition is computational, and thus appropriately conceptualized as transformations over representations. Nevertheless, embracing the metaphor of analog computation does have some striking "revolutionary" implications for cognitive science; here I'll briefly survey only three: representations are kinematic, simulations are ineliminable, and computability constrains functional analysis.

§4.1. *Representation: Kinematic and architectural, not static and symbolic*

Explanations of the MONIAC appeal to representations, and the complexity of the system implies that it cannot be adequately understood or explained without these representational ascriptions. Nevertheless, the vehicles of representation in the MONIAC are radically different from the vehicles of representation in a digital computer.

Digital computation is defined over an alphabet of discrete symbols. Each symbol represents only in virtue of some conventional interpretation, or assignment of content—'1' may be assigned the interpretation *one*, but it could just as easily be used to signify *zero*, the command *erase RAM*, or the person *Richard Nixon*. Symbols are static objects, and they suggest the characteristic transformations identified by Turing, namely they may be read or written to tape.

In contrast, representations in an analog system are typically kinematic, that is, movement or *change* in some quantity measures or represents *change* in some other quantity. No fixed amount of water in the MONIAC represents anything, but the kinematics of hydraulic flow do perform representational functions. Change in rate of flow represents change in rate of monetary circulation; change in height of surplus balances tank represents change in interest rate, etc.[9] These kinematic representations contrast with symbolic representations not only in the form of their vehicles, but also in the nature of their content: symbols most naturally represent objects or properties, while kinematic changes most naturally represent *quantities*, values that fall within a spectrum of possibilities, or scale. A moral for embodied cognition is to think of mental representations as analogous to measurements, as the theory of measurement provides the canonical account of the representation of quantities.[10]

A second form of representation in the MONIAC might be called architectural: aspects of its arrangement or physical structure represent mathematical relationships between kinematically represented quantities. The curvature of the wall of the surplus balances tank discussed above is one such structural feature. The value of interpreting these as representations of mathematical relations rather than merely instantiations of them is that these architectural features are subject to correctness conditions. If the outer wall of the surplus balances tank is bent over time, through mechanical wear and tear, then it will come to misrepresent Keynes' theory of the relationship between liquidity preference and interest rate, and thus computations with the MONIAC that assume Keynes' theory may be incorrect.

§4.2. *Explanation: Models, not modules*

One feature of dynamical models that has been stressed before, and which is vividly illustrated by the MONIAC, is their holistic character. While the classical digital metaphor suggested cognition be thought of as modular, or comprising informationally encapsulated sub-routines (Fodor, 1983), the dynamical, enactive approach suggests all features of the system are entangled in webs of mutual influence and interaction (Anderson, Richardson, and Chemero, 2012). One implication of holism is that simulation may play an ineliminable role in explanations of complex cognitive capacities

This holistic character of complex dynamical systems, and the corresponding explanatory importance of simulations, is nicely illustrated by one of the theoretical contributions of the MONIAC: deflation of the dispute between Keynes and Robertson over the determination of interest rates. Keynes argued that interest rates are determined by the preference for liquidity, i.e. for *holding* money rather than bonds, while Robertson argued that interest rate is determined by the demand for loanable funds, i.e. the differential "*flow* of saving versus that of investment" (Morgan, 2012, 209). In Phillips' words, participants in the debate "suffered through lack of a suitable technique for showing the process of change through time of the inter-related factors" (1950, 299). Essentially, Keynes and Robertson were each describing *part* of the economic dynamics in isolation, and this piecemeal approach obscured the elaborate interaction between these respective parts. The MONIAC, by instantiating both theories in a single system demonstrates they are "neither inconsistent with each

---

[9] We can see now why van Gelder could not find representations of the static sort in the Watt Governor. Considered as an analog computer, we should not expect any particular, fixed feature of the Governor, such as the angle of the arms, to do representational work. If any aspect of the system represents, it should do so kinematically—for instance change in the angle of the arms, or rate of change, are more appropriate candidate representations.

[10] For an extended discussion of how perceptual representations in particular are analogous to measurements, see Isaac, 2014.

other . . ., nor merely different ways of saying the same thing . . ., but are complementary parts of a wider system" (299). This dispute, and its resolution by the MONIAC, also nicely illustrates again why differential equations by themselves may not be enough to sufficiently explain a dynamical system. The further ingredient here is an understanding of the dynamics the separate equations jointly determine, and in the case of the interest rate debate this understanding was imparted through the explicit, visually accessible, ostension of those dynamics in the MONIAC.[11]

The constitutive role of models and simulations in the explanations of embodied cognitive science has been emphasized before, but suggested as further evidence against the explanatory value of representation attribution (e.g. Chemero, 2012, 99–101). A much-discussed example, for instance, is the HKB model of the synchronization of bilateral finger motion (Haken, Kelso, and Bunz, 1985). This model takes the transitions of finger wagging in and out of phase as a consequence of the nonlinear coupling of self-sustaining oscillators, reducing the change in phase ϕ to a function of two parameters, which may be fit to observed system dynamics:

$$V(\phi) = -a\cos\phi - b\cos 2\phi$$

No part of the model need be explained in representational terms, yet the close replication of the phase-shifting patterns observed in finger wagging experiments it produces lends it credence as an explanation. Oscillators coupled in accordance with this model characteristically exhibit spontaneous changes in relative phase, preceded by an increase in system entropy. Chemero argues that this general pattern is observed in, and thus extensions of this model explain, a variety of more complex phenomena, including synchronization across movements of other limbs, metrical phenomena in speech, and the "a-ha" moment that occurs when a subject switches to a new mode of problem solving (2012, 85–96).

Like the Watt Governor, however, the HKB model is extremely simple, and one might wonder whether non-representational models of this form can really "scale up" to account for complex cognitive behavior. The recommendation here is that scaling up will require much more elaborate models, which, though comprising qualitatively similar components, will typically demand a representational interpretation, on pain of failure to impart understanding. Nevertheless, *pace* traditional computationalism, mere identification of representational content and functional relationships between vehicles will likewise not, in general, be sufficient for an explanation—this follows from the inherently holistic, non-modular nature of cross-coupling within a dynamical system, even one performing a computation. Rather, the model must also be implemented in a simulation if its complex dynamics are to be fully understood. The general moral for embodied, enactive cognitive science, then, is that running of simulations and identifying representations are complementary, rather than exclusionary, steps in the construction of full-fledged explanations.

---

[11] As Morgan (2012) recounts, this point was crystal clear to the audience at Phillips' first, 1949, demonstration of the MONIAC at the LSE, while Robertson himself, having not seen the machine in person, seemed unable to grasp it from the paper alone (209).

§4.3. *Analog vs. Turing computability*

If a system is a computer, then it implements an effective procedure, i.e. one that requires only finite physical resources and robustly returns the same result on repetition. Computability theory analyzes effective procedures in terms of the (mathematical) functions they may compute. So, if cognition is analog computation, it consists only in functions (isomorphic to those) that are analog-computable, and thus the mathematical theory of analog computation should constrain and inform explanations in embodied cognitive science.

One project in the mathematics of analog computation investigates "effective" operations defined over mathematically continuous domains, e.g. the real line, and argues that these include non-Turing-computable functions, such as the halting problem, thus instantiating hyper-Turing computation (e.g. Siegelmann, 2003). More generally from this perspective, analogs to the key notions for Turing computability over integers (e.g. recursive functions, decidability, the complexity hierarchy[12]) may be defined for analog computation over reals (Blum, Shub, and Smale, 1989).

However, any real world analog device will be subject to systemic low-level noise—in the MONIAC, for instance, noise manifests as fluctuations in the movement of water at a scale smaller than the precision bound on valves and channels—, and thus procedures that are effective *in reality* cannot access the full precision of the reals. Consequently, although analog computers solve differential equations, and differential equations are typically defined over real numbers, the solutions delivered by realistic analog computers are at best approximations to real-valued solutions. This principled limit on the mathematical resolution of any particular analog computer is called by Vergis, Steiglitz, and Dickinson (1986) its "*absolute precision*, ε":

> [T]he solution obtained from [an analog computer] does not change when the physical variables range over an ε-neighborhood of their nominal values. . . . In order to discuss the operation of physical devices using mathematical models, it is important to insure that the models are robust in the sense that the physical behavior predicted by the mathematical model is not more sensitive to small changes in the model than is the underlying physical system to small perturbations. (94)

Consequently, if we are to investigate the properties of realistic analog computers mathematically, we need to include the bound on precision ε in our mathematical model of the device. Given a mathematical model of realistic analog computers such as this, one may ask whether they are computationally equivalent to realistic digital computers (i.e. those modeled by Turing machines with bounded tape length). If so, then it would constitute a further confirmation of the Church-Turing thesis and constrain cognitive explanations to only include Turing computable functions. Personally, I'm inclined to think this is the likely result, but the technical issues here

---

[12] See, e.g., Cutland, 1980, for the basics of Turing computability and equivalent models, and Li and Vitányi, 2003, for a thorough introduction to computational complexity theory.

are subtle and, as yet, not a matter of consensus. Nevertheless, we can still consider hypothetical implications of this line of research for embodied cognitive science.

Suppose it turns out that realistic analog computation is not only Turing equivalent, but that the coarse complexity classes of analog computation are equivalent to those of digital computation, i.e. problems that are NP-hard for digital computers (intuitively, those with no efficient, or "shortcut" solutions) may likewise only be solved inefficiently by analog computers; call this the *Complexity-Matching Thesis* (CMT). Vergis, Steiglitz, and Dickinson (1986) explicitly defend CMT, and it follows for the special case of cognition from Rubel's (1985) conjecture that brains utilize the all-or-none character of nervous activity to "digitally simulate" analog computations. A novel research program for embodied cognition could arise from combining the CMT and the *P-Cognition Thesis* (PCT, van Rooij, 2008), namely the posit that all cognitive operations are tractable, or computable in polynomial-time. For instance, if an enactivist model of some cognitive task is intractable, then CMT plus PCT implies the task analysis is incorrect, and the brain must actually be computing some simpler, more tractable function.[13] Conversely, and more interestingly, suppose it turns out that CMT is false, and that the complexity hierarchy of analog computations is not equivalent to the standard hierarchy. Then embracing PCT opens a new realm for the empirical comparison of embodied and classical models; for instance, if a traditional analysis of a putative cognitive task is (Turing) untractable, but an embodied analysis of the same task is (analog) tractable, the embodied analysis should be preferred.

Finally, even in the absence of specific results, computability considerations may substantively inform the interpretation of embodied models of cognition. Specifically, the constraint that computations be effective procedures implies that only dynamical behavior robust in the face of low-level noise is properly cognitive. Consider, for instance, the HKB model discussed above: does the phase-switching behavior of the model depend on arbitrarily small fluctuations in the system dynamics, or is it robust to low-level noise? If the former, then the behavior it models is not computational, and thus also not cognitive. More generally, for any putative instance of embodied cognition, one may ask whether it constitutes an effective procedure or not, and if not, it would seem to be an artifact of contingent features of the system dynamics, rather than a true cognitive capacity. This returns us to the consideration raised at the start of §2: if one assumes that cognitive capacities are adaptive, then they must also be robust to noise, i.e. constitute effective procedures. This follows from the observation that selective pressures cannot target one-off events; in order to be a target for natural selection a capacity or feature must be inheritable, and thus it must in some sense be robustly repeatable.

5. *Conclusion*

The Watt Governor is too simplistic a model to spawn convincing intuitions about embodied, enactive cognition. I suggest analog computers may serve as a better metaphor for multi-E cognitive science, and offer the MONIAC as a highly suggestive and easily visualizable exemplar. From this perspective, embodied, enactive

---

[13] For an extended survey and discussion of this method of reasoning in the context of traditional cognitive science, see Isaac, Szymanik, and Verbrugge, 2014.

cognitive science is consistent with computationalism, understood as the thesis that cognition comprises transformations over representations. Nevertheless, embracing an analog computation account of cognition has revolutionary consequences: representations should be understood as kinematic or architectural, explanations may constitutively involve simulations, and the mathematics of analog computability and complexity should constrain and inform the construction of embodied explanations.

Acknowledgements

These ideas were developed largely in response to conversations with Tony Chemero, Joe Dewhurst, Dan Hutto, Marcin Miłkowski, Paul Schweizer, Mario Villalobos, and Dave Ward; they have benefitted from the feedback of participants at the Edinburgh workshop on *Computational Psychology*, March 27, 2017, and the 9th Bethe Center workshop on *Computational Sciences and Reality*, October 6, 2017, where earlier versions were presented.

Bibliography

Anderson, M. L., M. J. Richardson, and A. Chemero (2012) "Eroding the Boundaries of Cognition: Implications of Embodiment," *Topics in Cognitive Science* 4: 1–14.

Barr, N. (2000) "The History of the Phillips Machine," in R. Leeson (ed.) *A. W. H. Phillips: Collected Works in Contemporary Perspective*. Cambridge UP. pp. 89–114.

Bissell, C. (2007) "The Moniac: A Hydromechanical Analog Computer of the 1950s," *IEEE Control Systems Magazine* 27: 69–74.

Blum, L., M. Shub, and S. Smale (1989) "On a Theory of Computation and Complexity over the Real Numbers: *NP*-Completeness, Recursive Functions and Universal Machines," *Bulletin (new series) of the American Mathematical Society* 21: 1–46.

Chemero, A. (2000) "Anti-Representationalism and the Dynamical Stance," *Philosophy of Science* 67: 625–647.

Chemero, A. (2009) *Radical Embodied Cognitive Science*. MIT Press.

Cutland, N. J. (1980) *Computability: An Introduction to Recursive Function Theory*. Cambridge UP.

Dewhurst, J. (2016) "Individuation without Representation," *British Journal for the Philosophy of Science*, axw018: https://doi.org/10.1093/bjps/axw018

Fodor, J. (1983) *The Modularity of Mind*. MIT Press.

Haken, H., J. A. S. Kelso, and H. Bunz (1985) "A Theoretical Model of Phase Transitions in Human Hand Movements," *Biological Cybernetics* 51: 347–356.

Hutto, D. D. and E. Myin (2013) *Radicalizing Enactivism*. MIT Press.

Isaac, A. M. C. (2014) "Structural Realism for Secondary Qualities," *Erkenntnis* 79: 481–510.

Isaac, A. M. C. (*forthcoming*) "Review of G. Piccinini *Physical Computation: A Mechanistic Account*," *Philosophical Review*.

Isaac, A. M. C., J. Szymanik, and R. Verbrugge (2014) "Logic and Complexity in Cognitive Science," in A. Baltag and S. Smets (eds.) *Johan van Benthem on Logic and Information Dynamics*, Springer. pp. 787–824.

Li, M. and P. Vitányi (2008) *An Introduction to Kolmogorov Complexity and Its Applications*, 3rd ed. Springer.

Miłkowski, M. (2015) "The Hard Problem of Content: Solved (Long Ago)," *Studies in Logic, Grammar, and Rhetoric* 41(54): 73–88.

Morgan, M. S. (2012) *The World in the Model*. Cambridge UP.

Phillips, A. W. (1950) "Mechanical Models in Economic Dynamics," *Economica* 17(67): 283–305.

Piccinini, G. (2008) "Computation without Representation," *Philosophical Studies* 137: 205–241.

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

Rowlands, M. (2009) "Enactivism and the Extended Mind," *Topoi* 28: 53–62.

Rubel, L. A. (1985) "The Brain as Analog Computer," *Journal of Theoretical Neurobiology* 4: 73–81.

Schweizer, P. (2017) "Cognitive Computation *sans* Representation," in T. Powers (ed.) *Philosophy and Computing: Essays in Epistemology, Philosophy of Mind, Logic, and Ethics*, Philosophical Studies Series, vol. 128. *Springer*. pp. 65–84.

Shagrir, O. (2010) "Brains as Analog-Model Computers," *Studies in History and Philosophy of Science* 41: 271–279.

Shannon, C. E. (1941) "Mathematical Theory of the Differential Analyzer," *Studies in Applied Mathematics*, 20: 337–354.

Shapiro, L. (2010) *Embodied Cognition*. Routledge.

Siegelmann, H. T. (2003) "Neural and Super-Turing Computing," *Minds and Machines* 13: 103–114.

Soroka, W. W. (1954) *Analog Methods in Computation and Simulation*. McGraw-Hill Publishing Co., Ltd.

Turing, A. M. (1936) "On Computable Numbers with an Application to the Entscheidungsproblem," *Proceedings of the London Mathematical Society* 2-42(1): 230–265.

van Gelder, T. (1995) "What Might Cognition Be, If Not Computation?" *Journal of Philosophy* 92: 345–381.

van Rooij, I. (2008) "The Tractable Cognition Thesis," *Cognitive Science* 32: 939–984.

Varela, F. J., E. Thompson, and E. Rosch (1991) *The Embodied Mind*. MIT Press.

Vergis, A., K. Steiglitz, and B. Dickinson (1986) "The Complexity of Analog Computation," *Mathematics and Computers in Simulation* 28: 91–113.

Villalobos, M. and J. Dewhurst (2017) "Enactive Autonomy in Computational Systems," *Synthese*: https://doi.org/10.1007/s11229-017-1386-z

D. Ward, D. Silverman, and M. Villalobos (2017) "Introduction: The Varieties of Enactivism," *Topoi* 36: 365–375.

Wheeler M. (2017) "The Revolution will not be Optimised: Radical Enactivism, Extended Functionalism, and the Extensive Mind," *Topoi* 36: 457–472.