

# Conscious Enactive Computation

Daniel Estrada

New Jersey Institute of Technology, Newark NJ 07102  
djestrada@gmail.com

**Abstract.** This paper looks at recent debates in the enactivist literature on computation and consciousness in order to assess conceptual obstacles to building artificial conscious agents. We consider a proposal from Villalobos and Dewhurst (2018) for enactive computation on the basis of organizational closure. We attempt to improve the argument by reflecting on the closed paths through state space taken by finite state automata. This motivates a defense against Clark’s recent criticisms of “extended consciousness”, and perhaps a new perspective on living with machines.

**Keywords:** enactivism, artificial intelligence, computation, Turing machine, state space, finite state automata, predictive coding, consciousness

## 1 Introduction

Enactivism challenges the dominant cognitive paradigm in psychology with an account of intentional (purposive) agency that is grounded in the emergent dynamics of biological complexity [15, 43, 46]. Specifically, enactivism holds that biological life is characterized by *adaptive self-constitution*: living systems construct and maintain their own organized structure through their active engagement with a changing world [4, 35]. This approach motivates a systematic account of autonomy [3, 33, 41, 48], intentional agency [17, 31], subjective consciousness [19, 28], and identity in complex dynamical systems [5, 6], with the promise of a consistent and unified explanatory framework across the full range of biological processes, from the biomechanics of single-celled organisms to ecologies and societies [18, 26, 44].

Despite the emphasis on biological complexity, enactivism has from its inception maintained a robust research program investigating artificial intelligence, artificial life, and robotics (hereafter AI) [1, 2, 13, 16, 20, 42]. This research aims to develop models, simulations, and robots that assist in the scientific investigation of biological complexity and adaptive systems. For instance, AI that exhibits some dynamically self-organizing behavior might serve as a useful “proof of concept” demonstrating key enactivist principles (see [20] for examples). However, while robotics research has already felt a significant impact from the embodied approach [37, 38], enactivist AI is often advanced against a backdrop of criticism directed at “merely” computational or representational explanations [22, 23]. As a founder of enactivism Francisco Varela put it, “This fundamental paradigm of the digital computer program will not do for biology, nor for AI.” [46]

A recent set of papers from Villalobos, Dewhurst, Ward and colleagues (hereafter Villalobos) [14, 49–51] address these historical tensions between enactivism and computation. Villalobos argues that the enactivists are mistaken to treat computers as mere symbolic processors of abstract representations. Drawing on a mechanist account of computation, Villalobos suggests an interpretation of the classical Turing machine which they claim would meet enactivist conditions for *self-determination*. If so, it would suggest that embodied agency could be given a computational rather than biological basis without sacrificing enactivism’s theoretical commitments to the dynamical interactions between agent and world. This argument strikes at the foundations of the enactivist program, and threatens to overturn more than twenty years of enactivist thought on AI and computation.

The central concern of this paper is to assess the proposal for enactive computation put forward by Villalobos. Their argument turns on the enactivist interpretation of self-determination in terms of *organizational closure*. While we think Villalobos’ examples fail to meet strong enactivist conditions on closure, we suggest they can be improved through explicit consideration of the structure of the finite state automata (**FSM**) that controls a classic Turing machine. This highlights an important form of closure that is, we argue, more fundamental than organizational closure: namely, the closed path through state space taken by the FSM. We claim that computation is fundamentally concerned with the structure of paths through state space, and that all living organisms can be characterized by such paths. This result suggests computation as the fundamental basis from which the enactivist program must emerge. We then consider the implications of this argument for a particular strand of criticism raised by Clark [10, 12] against enactivist proposals for “extended consciousness” [36]. We conclude with general thoughts on the implications these arguments have for living with machines.

## 2 Organizational closure and Turing’s machine

Organizational closure serves as the basis for the enactivist approach to autonomous intentional (purposive) behavior, and names the sense in which biological organisms are *self-determined* [3, 4, 47]. A system is *organized* when its constitutive components are arranged into a network of functionally interdependent processes and constraints [27]. An organization is *closed* when the operation of its constitutive components are themselves sufficient for the adaptive construction and generation of its organized state [35]. Enactivists argue that organizational closure provides an *intrinsic* basis for identifying organisms and their boundaries as unified wholes. Furthermore, enactivists emphasize that organisms are precariously situated within a dynamic world to which they must continually adapt in order to maintain their constitutive organization. This *precariousness* creates conditions that demand coordinated action from the organism as a unified system [7]. This gives rise to what enactivists call *adaptive sense-making*, which serves as the basis for investigations into consciousness and phenomenology [19, 28, 43].

Beyond its central role in the enactivist theory of autonomous agency, organizational closure also figures in enactivist criticisms of classical computation<sup>1</sup>. Enactivists contrast the closed structure of biological organisms with the *open* or *linear* structure of traditional computing machines [20]. On this view, computers operate through a sequence of formal operations that transforms symbolic “input” into symbolic “output”. Enactivists claim at least two important differences between computation and the adaptive self-constitution of biological organisms. First, computers perform stepwise formal operations on symbolic input, rather than performing dynamic mechanical operations within a changing world. Second, computers don’t “build themselves” in the sense relevant for adaptive self-constitution, which requires organizational closure. Put simply, computers aren’t self-determined wholes with a world of their own, and so cannot serve as the intrinsic subject of an experience. Instead, computers are artifacts created through external processes of human design and manufacturing. Such considerations lead Froese and Ziemke [20] to distinguish the *behavioral autonomy* characteristic of certain kinds of self-controlled machines (say, a dishwasher on a timer), from the *constitutive autonomy* characteristic of living biological systems.

Villalobos’ argument for enactive computation in [51] is designed to show that a Turing machine can meet the conditions for self-determination as described by Maturana (1988) [30]. Here, self-determination is identified with *functional closure*. A system has functional closure when its organizational structure contains closed feedback loops. As an example, Villalobos offers a thermostat regulating the temperature of a house. The behavior of the thermostat-house system is characterized by a feedback loop between these two components which satisfies functional closure on Maturana’s definition. Of course, while the thermostat-house system “controls itself” with respect to temperature, it is not adaptively self-constituting in any deeper sense; thermostats and houses don’t build themselves with their parts alone. Thus, functional closure is not sufficient for organizational closure of the sort required for constitutive autonomy. Nevertheless, Villalobos argues this control structure does not connect inputs to outputs through a linear sequence of symbolic processes, and so is not “open”. It is, they argue, closed and minimally self-determining in a sense relevant for enactivist theory.

Villalobos then applies this feedback loop model to the classic Turing machine. Turing [45] proposed a computing machine with three components: a tape with discrete cells; a read-write head that operates on the tape; and a program which controls the operation of the head. On the enactivist interpretation, the tape serves input to the machine and records output from the machine, and the machine (the head and program) performs formal operations that convert the former to the latter as a linear process. Against this view Villalobos offer an alternative, inspired by Wells [54] and Piccinini [39, 40], that interprets the

---

<sup>1</sup> Enactivists are not universally hostile to computation. Importantly, Mossio et al [34] render an organizationally closed system in the  $\lambda$ -calculus, and argue that “there are no conceptual or principled problems in realizing a computer simulation or model of closure.” Such arguments have resulted in a split between radical anti-computationalists [22] and more traditional versions of enactivism. See [8, 53].

Turing machine in terms of looping interactions between the machine and the tape. This forms a functionally closed loop, much like the thermostat-house system, which implies self-determination in the sense that the computer’s state is determined by the interactions between the machine and the tape. In an analog computer these constraints might appear as features of the physical mechanisms of the device, thereby eliminating any symbolic aspect of the computation. Thus, Villalobos argues, even a classical Turing machine can be understood as purely mechanical and functionally closed, and so evades both enactivist criticisms of computation. While this argument doesn’t entail that computers are conscious living creatures of equivalent complexity to biological organisms, it does confront a major hurdle within the enactivist literature to treating computing machines as genuinely purposive agents with a world of their own.

Does Villalobos’ argument succeed? Briefly, no: functional closure alone is not sufficient for adaptive self-constitution of the sort relevant for intentional agency or adaptive sense-making. Villalobos’ ‘enactive’ Turing machine is merely behaviorally and not constitutively autonomous. While Maturana’s account is influential, recent work has developed more rigorous constraints on organizational closure. For instance, Mossio et al. [32,35] present a model of closure which requires that constitutive constraints operate across multiple scales or levels of organization to achieve closure. While the thermostat-house system is functionally closed, we might say that closure occurs at a single scale, namely the feedback loop that controls temperature. At other scales, for instance the internal structure of the thermostat mechanism, the system is not closed or self-determining but depends directly on external processes. Similarly, Turing’s machine appears to be functionally closed only at the level of operations of the head on the tape and nowhere else. Biological systems, on the other hand, are in some sense self-determining *all the way through*—or at least they are self-organized across a range of scales from inter-cellular biochemistry through geopolitics that covers the breadth of our experiences of a meaningful world as human agents. A Turing machine might be functionally closed, but it covers nothing close to the same range of interactivity.

How many levels of organizational constraints are required to distinguish between behavioral and constitutive autonomy? Mossio’s model suggests at least two. If so, Villalobos’ argument might be improved by describing a Turing machine with two layers of self-determining organizational constraints rather than one. In the next section, I will discuss how the classic Turing machine already captures organizational closure across two layers of constraint.

### 3 Closed paths through state space

If we suspend the anti-representational commitments of enactivism for a moment, there’s an important feature of Turing’s machine which is not explicitly addressed in these arguments: the structure of the program which controls the read-write head. In Turing’s model, the program takes the form of a finite state machine (**FSM**). FSMs are abstract automata that are characterized by a finite

number of discrete states, a set of rules that describe the conditions for transitioning between states depending on what is read from the tape. These rules can be represented as a state transition table, which can be realized<sup>2</sup> in a physical machine in a number of ways. The physical Turing machine is ‘programmed’ insofar as it realizes the abstract state transition structure of the FSM.

The abstract nature of the FSM should not worry enactivists [27]. An FSM can in principle be realized by simple physical mechanisms; there’s nothing inherently “symbolic” about the FSM. The FSM is not necessarily used by a computer to “represent the world”. The FSM is just an abstract model of the states a machine can be in, and the conditions for transitioning between these states. Enactivist literature is often directly preoccupied with systems being in certain states, like the equilibrium state (homeostasis), and with the activities organisms must perform to maintain these states [25, 41]. To this extent, enactivist theory depends on state abstractions of the same sort used to describe the FSM. Describing the autonomy of an organism in terms of “organizational closure” is already to appeal to control structure that achieve an abstract state, so there should be no principled objections from enactivists to discussing the equally abstract structure of the FSM.

While the FSM can be represented as a transition table, it is also customary to represent an FSM with a *state space diagram* with states represented as circles, and arrows between circles representing the transitions between states. A state space diagram has a *closed path* (or *loop*) if some sequence of operations will return the system to a previous state. Such closed paths are typical in discrete, finite computing automata, but are also familiar from continuous cases in the physical world. Suppose I take water at room temperature, freeze it to ice, then let it thaw back to room temperature. The water crossed a state transition, then crossed back; we can represent this as a path through the state space of water that loops back to where it began and in this sense is closed across the relevant state transition. Homeostasis is an interesting state for living biological organisms precisely because they maintain the state as a *fixed point attractor*, returning to equilibrium after minor disturbances. This is another way of saying that homeostasis is characterized by a closed path in state space (**CPSS**).

With these considerations in mind, we propose that CPSSs, and paths in state space generally, are of fundamental relevance to enactivist models of self-determination. Moreover, CPSSs put computers and organisms on equal ontological footing. Recall the theoretical motivation for appealing to organizational closure to explain autonomy: it provides an intrinsic basis for individuating a system as a unified whole, and so serves as a basis for adaptive sense-making.

---

<sup>2</sup> For historical reasons originating with Putnam [21], it is often taken for granted that a definition of computation in terms of finite state automata cannot distinguish between different *realizations* of a computer, and so cannot in principle provide an explanation for cognitive behavior. Piccinini [39] cites this as an explicit motivation for developing his mechanistic account of computation. There are good reasons for thinking that Putnam’s concerns are overstated [9, 24], but this issue is beyond the scope of this paper. Thanks to Jon Lawhead for pointing this out.

We claim that a CPSS accomplishes the same theoretical task: organisms can be identified intrinsically as the collection of processes and constraints that walk a CPSS. This definition is intrinsic in the same sense as organizational closure: whether a path counts as “closed” is set by the constitution of the system itself and the state space it traverses. The abstraction of state space traversals is general enough to apply consistently across physics, biology, and computer science. More strongly, we claim that any organizationally closed system can be characterized by a collections of CPSSs with a fixed attractor at the constitutive organized state. This suggests that CPSSs are theoretically a more fundamental form of closure than organizational closure. Indeed, the important sense of ‘closure’ captured by the enactivists has less to do with daisy-chained functions looping on themselves, and more to do with the structure of the state space traversals those functional relationships enable. Strictly speaking, neither functional nor organizational closure is necessary for walking a CPSS.

Not every Turing machine will walk a CPSS, but it is exceedingly common for them to do so<sup>3</sup>. We can think of the CPSSs which characterize a Turing machine’s program as another scale of closure, one which directly controls the looping interactions between head and tape. With two scales of closed loops, this would appear to meet Mossio’s stronger constraints on closure, and thus we have shown the classical Turing machine might already constitute an adaptively self-constituting system on enactivist grounds. Or, perhaps more realistically, the depth of closure matters less than what states those functional relationships (closed otherwise) make available for the organism as it walks paths in state space.

## 4 Extended consciousness

To appreciate how CPSSs can be useful to enactivism, consider a recent debate on the bounds of consciousness. Despite his strong influence on enactivism, Clark has pushed back against attempts to locate processes constitutive of conscious experience in the world [10]. Clark argues there is no good reason to do so; the activity constitutive of a conscious experience occurs immediately within patterns of neural firings. Clark advocates for an explanatory approach called “predictive coding” which uses “a hierarchical generative model that aims to minimize prediction error within a bidirectional cascade of cortical processing” [12]. Clark argues that the model works by rapidly updating on the basis of new information. This leaves little bandwidth for external changes to impact the updating model beyond sensory input; the dominant influence on most neurons is the activity of other neurons. Thus, Clark argues, it is unlikely that external processes play a constitutive role in conscious experience.

Ward [52] offers a response to Clark on behalf of enactivists that appeals to multiple layers of interactions between the agent and world. Clark’s mistake, on

---

<sup>3</sup> The question of deciding in general whether a path in state space will close is formally equivalent to the halting problem, and so is not computable. See [29].

this view, is to localize consciousness to any single process in the organized hierarchy. The appeal to multiple layers should by now be a familiar enactivist move, one Clark rejects as superfluous in this case [11]. Whatever world-involving processes enactivists believe are important, Clark claims he can account for them with predictive coding that occurs strictly within neural activity. So consciousness appears stuck in the head.

Clark's alternative doesn't appeal to enactivists because the world-involving aspects of predictive coding appear linear and open, like a computer, rather than closed like an organism. This isn't an accurate perception; the cascade of neural activity develops with looping feedback until the neurons reach stability, so there *are* functionally closed processes; those processes just aren't extended and world-involving beyond the brain. Enactivists are attracted to externalism because they view consciousness as inherently world-involving and organizationally closed. Just as with Villalobos' computer, enactivists are hoping to find a closed organizational structure associated with the embodied conscious state. Since closure is an indicator of unification and wholeness, enactivists expect neural activity and world-involving processes to demonstrate dynamic functional interdependencies. Clark's argument that the neural activity is not functionally dependent on external processes is therefore fatal to extended consciousness.

Perhaps CPSSs can help resolve this conflict amicably? If we think about closure in terms of CPSSs we can recover the looping interactions that are inherently world-involving and *closed in state space*, while conceding to Clark that the neural activity is sufficiently explanatory of the neurophysiological interactions that give rise to the conscious state. By considering consciousness from the perspective of traversals in state space, we are no longer confined to a single closed loop of spanning organizational levels of physical or biological interactions. Instead, dynamical activity across different scales will form many different kinds of closed paths in different state spaces. Some of these CPSSs will be characterized by inherently world-involving *states*, and in this sense will recover an enactivist sense of extended consciousness compatible with predictive coding.

Consider, for instance, that it is easier to maintain your balance with your eyes open than closed. Here we have two cortical cascades that traverse through a conscious state: one producing visual experiences, and one producing motor activity to maintain balance. These two systems work in concert and reinforce each other. On the enactivist framing, maintaining balance is a precarious state that inherently involves the configuration of the body as a massive physical object with specific dimensions. Thus, the configuration of my body is a fundamental factor in whether I am in a balanced state<sup>4</sup>. From the perspective of state space, the balanced state is a fixed attractor for certain CPSSs; the neural cascades that produce my balancing behavior are associated with attempts to close a traversal in state space and return to the balanced state. This brings in looping, inherently world-involving processes into an explanation of my behavior as an

---

<sup>4</sup> This is the case whether or not I am explicitly conscious of my body's configuration as a physical object. I might be aware that I am balanced, but lack proprioceptive awareness of the bodily configuration that produces the balanced state.

agent without committing to implausible functional interdependencies between neurons and world. The important dependencies for closure, and ultimately for autonomy, identity, and consciousness, are found in state space.

## 5 Conclusion

We don't view CPSSs as a threat to enactivism's positive theory of autonomy or adaptive sense-making. Instead, we see it correcting the over-emphasized anti-computationalism that has historically motivated the view. We think enough speaks in favor of the enactive approach that it needn't appeal to a problematic ontological distinction between computing machines and biological life. Insofar as Villalobos' argument also serves these goals, this paper is meant to push harder in the same direction.

We believe the significance of this correction extends beyond enactivist theory, and is relevant to general debates concerning the relationship between computing machines and biological organisms. State space descriptions provide a convenient abstraction within which these distinctions cannot be motivated on a priori or ontological grounds, but must be defended by appeal to specific patterns of dynamical behavior. State space descriptions may also help us model the elaborate functional interdependencies between ourselves and the technosocial superstructures we inhabit, and thereby help us to better appreciate the capacities and perspectives of the machines we live alongside.

## References

1. Agmon, E., Egbert, M., Virgo, N.: The biological foundations of enactivism: A report on a workshop held at Artificial Life XV. *Artificial life* **24**(1), 49–55 (2018)
2. Aguilar, W., Santamara-Bonfil, G., Froese, T., Gershenson, C.: The Past, Present, and Future of Artificial Life. *Frontiers in Robotics and AI* **1** (Oct 2014)
3. Barandiaran, X.E.: Autonomy and enactivism: Towards a theory of sensorimotor autonomous agency. *Topoi* **36**(3), 409–430 (2017)
4. Bechtel, W.: Biological mechanisms: Organized to maintain autonomy. In: *Systems biology*, pp. 269–302. Elsevier (2007)
5. Bechtel, W.: Identity, reduction, and conserved mechanisms: Perspectives from circadian rhythm research (2012)
6. Bechtel, W.: Systems biology: Negotiating between holism and reductionism. In: *Philosophy of Systems Biology*, pp. 25–36. Springer (2017)
7. Burge, T.: Primitive agency and natural norms. *Philosophy and Phenomenological Research* **79**(2), 251–278 (2009)
8. Cárdenas, M.L., Letelier, J.C., Gutierrez, C., Cornish-Bowden, A., Soto-Andrade, J.: Closure to efficient causation, computability and artificial life. *Journal of Theoretical Biology* **263**(1), 79–92 (2010)
9. Chalmers, D.J.: Does a rock implement every finite-state automaton? *Synthese* **108**(3), 309–333 (1996)
10. Clark, A.: Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head. *Mind* **118**(472), 963–993 (Oct 2009)



11. Clark, A.: Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience. *Mind* **121**(483), 753–771 (2012)
12. Clark, A.: Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences* **36**(3), 181–204 (2013)
13. De Loor, P., Manach, K., Tisseau, J.: Enaction-based artificial intelligence: Toward co-evolution with humans in the loop. *Minds and Machines* **19**(3), 319–343 (2009)
14. Dewhurst, J., Villalobos, M.: The Enactive Automaton as a Computing Mechanism: Enactive Automaton as a Computing Mechanism. *Thought: A Journal of Philosophy* **6**(3), 185–192 (Sep 2017)
15. Di Paolo, E., Thompson, E.: The enactive approach. *The Routledge handbook of embodied cognition* pp. 68–78 (2014)
16. Di Paolo, E., Buhrmann, T., Barandiaran, X.: *Sensorimotor life: An enactive proposal*. Oxford University Press (2017)
17. Di Paolo, E.A.: Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the cognitive sciences* **4**(4), 429–452 (2005)
18. Froese, T., Di Paolo, E.A.: The enactive approach: Theoretical sketches from cell to society. *Pragmatics & Cognition* **19**(1), 1–36 (2011)
19. Froese, T., Gallagher, S.: Phenomenology and artificial life: toward a technological supplementation of phenomenological methodology. *Husserl Studies* **26**(2), 83–106 (2010)
20. Froese, T., Ziemke, T.: Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence* **173**(3-4), 466–500 (2009)
21. Hilary, P.: *Representation and reality*. Cambridge, Mass.: A Bradford Book (1988)
22. Hutto, D.D., Myin, E.: *Radicalizing enactivism: Basic minds without content*. MIT Press (2012)
23. Hutto, D.D., Myin, E.: *Evolving enactivism: Basic minds meet content*. MIT Press (2017)
24. Joslin, D.: Real realization: Dennetts real patterns versus putnams ubiquitous automata. *Minds and Machines* **16**(1), 29–41 (2006)
25. Kauffman, S.A.: *Investigations*. Oxford University Press (2000)
26. Kirchhoff, M.D., Froese, T.: Where there is life there is mind: In support of a strong life-mind continuity thesis. *Entropy* **19**(4), 169 (2017)
27. Levy, A., Bechtel, W.: Abstraction and the organization of mechanisms. *Philosophy of science* **80**(2), 241–261 (2013)
28. Lutz, A., Thompson, E.: Neurophenomenology integrating subjective experience and brain dynamics in the neuroscience of consciousness. *Journal of consciousness studies* **10**(9-10), 31–52 (2003)
29. Luz Crdenas, M., Letelier, J.C., Gutierrez, C., Cornish-Bowden, A., Soto-Andrade, J.: Closure to efficient causation, computability and artificial life. *Journal of Theoretical Biology* **263**(1), 79–92 (Mar 2010)
30. Maturana, H.: Ontology of observing: The biological foundations of self consciousness and the physical domain of existence. *The Irish Journal of Psychology* **9**(1), 25–82 (1988)
31. Merritt, M.: Thinking-is-moving: dance, agency, and a radically enactive mind. *Phenomenology and the Cognitive Sciences* **14**(1), 95–110 (2015)
32. Montvil, M., Mossio, M.: Biological organisation as closure of constraints. *Journal of Theoretical Biology* **372**, 179–191 (2015)
33. Moreno, A., Mossio, M.: *Biological autonomy*. A Philo (2015)
34. Mossio, M., Longo, G., Stewart, J.: A computable expression of closure to efficient causation. *Journal of Theoretical Biology* **257**(3), 489–498 (2009)

35. Mossio, M., Moreno, A.: Organisational closure in biological organisms. *History and philosophy of the life sciences* pp. 269–288 (2010)
36. Noë, A., Thompson, E.: Are there neural correlates of consciousness? *Journal of Consciousness studies* **11**(1), 3–28 (2004)
37. Pfeifer, R., Gomez, G.: Interacting with the real world: design principles for intelligent systems. *Artificial life and Robotics* **9**(1), 1–6 (2005)
38. Pfeifer, R., Iida, F., Bongard, J.: New robotics: Design principles for intelligent systems. *Artificial life* **11**(1-2), 99–120 (2005)
39. Piccinini, G.: Computing mechanisms. *Philosophy of Science* **74**(4), 501–526 (2007)
40. Piccinini, G.: *Physical computation: A mechanistic account*. OUP Oxford (2015)
41. Ruiz-Mirazo, K., Moreno, A.: Basic autonomy as a fundamental step in the synthesis of life. *Artificial life* **10**(3), 235–259 (2004)
42. Suzuki, M., Floreano, D.: Enactive robot vision. *Adaptive Behavior* **16**(2-3), 122–128 (2008)
43. Thompson, E.: Life and mind: From autopoiesis to neurophenomenology. A tribute to Francisco Varela. *Phenomenology and the cognitive Sciences* **3**(4), 381–398 (2004)
44. Thompson, E.: *Mind in life: Biology, phenomenology, and the sciences of mind*. Harvard University Press (2010)
45. Turing, A.M.: On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London mathematical society* **2**(1), 230–265 (1937)
46. Varela, F., Thompson, E., Rosch, E.: *The Embodied Mind: Cognitive Science and Human Experience* MIT Press. Cambridge, Massachusetts (1991)
47. Varela, F.G., Maturana, H.R., Uribe, R.: Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems* **5**(4), 187–196 (1974)
48. Vernon, D.: Enaction as a conceptual framework for developmental cognitive robotics. *Paladyn, Journal of Behavioral Robotics* **1**(2), 89–98 (2010)
49. Villalobos, M., Dewhurst, J.: Computationalism, Enactivism, and Cognition: Turing Machines as Functionally Closed Systems. In: *AIC*. pp. 138–147 (2016)
50. Villalobos, M., Dewhurst, J.: Why post-cognitivism does not (necessarily) entail anti-computationalism. *Adaptive Behavior* **25**(3), 117–128 (2017)
51. Villalobos, M., Dewhurst, J.: Enactive autonomy in computational systems. *Synthese* **195**(5), 1891–1908 (May 2018)
52. Ward, D.: Enjoying the Spread: Conscious Externalism Reconsidered. *Mind* **121**(483), 731–751 (Jul 2012)
53. Ward, D., Silverman, D., Villalobos, M.: Introduction: The varieties of enactivism. *Topoi* **36**(3), 365–375 (2017)
54. Wells, A.: Turing’s analysis of computation and theories of cognitive architecture. *Cognitive Science* **22**(3), 269–294 (1998)