

Resolving Two Tensions in 4E Cognition Using Wide Computationalism

Luke Kersten (L.M.Kersten@sms.ed.ac.uk)

School of Philosophy, Psychology, and Language Sciences, University of Edinburgh
3 Charles Street, EH89AD, Edinburgh, UK

Joe Dewhurst (s0819473@sms.ed.ac.uk)

School of Philosophy, Psychology, and Language Sciences, University of Edinburgh
3 Charles Street, EH89AD, Edinburgh, UK

George Deane (s1670299@sms.ed.ac.uk)

School of Philosophy, Psychology, and Language Sciences, University of Edinburgh
3 Charles Street, EH89AD, Edinburgh, UK

Abstract

Recently, some authors have begun to raise questions about the potential unity of 4E (enactive, embedded, embodied, extended) cognition as a distinct research programme within cognitive science. Two tensions, in particular, have been raised: (i) that the body-centric claims embodied cognition militate against the distributed tendencies of extended cognition and (ii) that the body/environment distinction emphasized by enactivism stands in tension with the world-spanning claims of extended cognition. The goal of this paper is to resolve tensions (i) and (ii). The proposal is that a form of ‘wide computationalism’ can be used to reconcile the two tensions and, in so doing, articulate a common theoretical core for 4E cognition.

Keywords: 4E cognition, wide computationalism, body-centrism, extended functionalism, autopoietic theory

Introduction

Enactive, embodied, embedded and extended cognition, or simply 4E cognition, has often been thought to form a collective challenge to traditional or classical cognitive science (Menary, 2010). Common to many of these views is the idea that cognitive processes are often integrated with and heavily dependent on bodily and environmental structures (Varela, Thompson, & Rosch, 1991; Clark & Chalmers, 1998; Haugeland, 1998; Hutto & Myin, 2013).

More recently, some authors have begun to raise questions about the potential unity of 4E cognition as a distinct research programme within cognitive science (Clark, 2008a; Clark & Kiverstein, 2009; Menary, 2010). Two tensions, in particular, have been raised: (i) that the body-centric claims embodied cognition militate against the distributed tendencies of extended cognition and (ii) that the body/environment distinction emphasized by enactivism stands in tension with the world-spanning claims of extended cognition. These two tensions constitute a problem for cognitive science insofar as 4E cognition is thought to form distinct field of study, and not merely a loose set of alphabetically related approaches (Ward & Stapleton, 2012).

The goal of this paper is to resolve tensions (i) and (ii). The proposal is that a form of ‘wide computationalism’ can be used to reconcile the two tensions, and, in so doing,

articulate a common theoretical core for 4E cognition. It is argued that wide computationalism satisfies the various demands of the embodied, enactive and extended theorists in virtue of placing a simultaneous emphasis on abstract analysis and functional mechanisms.

Two Tensions in 4E Cognition

Following Clark and Kiverstein (2009), three ‘strands’ can be identified as generating the two tensions within 4E cognition, these include: body-centrism, extended functionalism, and autopoietic theory.

The first strand is body-centrism. This is the idea that the body has a non-trivial role in determining mental states and functioning, that the details of a creature’s embodiment have a profound affect on the nature and functioning of the mind (Noë & Reagan, 2001; Noë, 2004; Gallagher, 2005). Shapiro, for instance, writes: “[P]sychological processes are incomplete without the body’s contributions. Vision for human beings is a process that includes features of the human body. This means that a description of various perceptual capacities cannot maintain body-neutrality” (2004, p.190). The body is depicted as ‘intrinsically special.’ Body-centrists hold that without discussion of the unique contribution of bodily structures and activities, cognitive explanations are crucially lacking. The view is also sometimes called the “constitutive-contribution claim” (Clark, 2008a).

Support for body-centrism comes from research highlighting the functional dependences of mental processes on bodily structures and activities (Clark, 2008a, b). Work on embodiment and conceptualization, for example, demonstrates that understanding abstract concepts, such as love, often depends on the metaphorical expansions of more familiar concepts, such as up and down or front and back (Lakoff & Johnson, 1980).

The second strand is extended functionalism. For the extended functionalist, cognitive systems are functional wholes distributed across diverse sets of components and processes. Cognitive activities involve a complex balancing act between brain, body and world (Harman 1998, Clark & Chalmers, 1998; Wilson 2004; Wheeler, 2010). The spirit of

extended functionalism is embodied in what Clark and Chalmer's (1998) call the 'parity principle', which says: "[i]f, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as a part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process" (p.29). The parity principle stresses the location neutrality of cognitive analysis. It highlights abstract, functional analysis in place of detailed physiological investigations. The view is also sometimes called the "distributor role" in discussions of embodied cognition (Wilson & Foglia, 2016).

Support for extended functionalism comes from research focusing on the way in which cognizers often exploit, scaffold, and distribute cognitive activities across bodily and environment structures (Clark, 2005, 2008b). Work on problem solving, for example, shows that people often simplify and transform complex problems, such as the Tower of Hanoi, by manipulating physical environments (Kirsh & Maglio, 1995; Kirsh, 2009).

The final strand to consider is 'autopoietic theory.' The central claim of autopoietic theory is that cognitive systems are created by the reciprocal interaction of internal and external components in the service of some larger function, such as homeostasis (Weber and Varela, 2002). Autopoietic theory connects to 4E cognition via the notion of 'sense making'.

Autopoietic theory maintains that because living systems, such as cognition, are autonomous, self-regulating systems, and sense making is required for maintaining a system's boundary, autopoietic systems produce and maintain a physical boundary between the organism and its physical environment. Because sense making is a self-regulating act, organisms often bring forth meaning on the basis of their autonomy – autonomy in this context means actively sustaining identity under precarious circumstances. A system maintains its organization by regulating its interactions with the environment via sense making (Thompson, 2007; Wheeler, 2009).

Consider how each of the three strands fit within 4E cognition. First, enactivism and embodied cognition often endorse body-centrism in virtue of emphasizing what they take to be the unique contributions of bodily structures and activities, such as sensorimotor knowledge. Second, extended cognition often endorses extended functionalism by assigning a non-trivial role to environmental elements in sustaining cognitive activities. Third, enactivism is often framed in terms of autopoietic theory insofar as sense making is treated as a constitutive element of demarcating the organism/environment boundary.

Not every version of enactivism is committed to autopoietic theory, and not every version of embodied cognition is committed to body-centrism. There is, at least in principle, some compatibility between the various views. Nonetheless, because some versions of each view are, as a matter of fact, committed to the different strands, the two

tensions do represent a substantial challenge for 4E cognition.

Consider, then, how the three strands generate the two tensions. The first tension follows from the fact that if the body has a non-trivial role in determining mental states, then cognition cannot also be location neutral; the converse of which is that if cognition is location neutral, then the body cannot have a privileged status in cognition. If the body is simply an instrument through which larger functional complexes are realized, then bodily structures cannot form the exclusive realization base of cognitive activities. Extended functionalism precludes the constitutive contribution claims of body-centrism, while body-centrism precludes the possibility of cognitive systems extending beyond the boundary of the individual.

The second tension emerges from the idea that if extended functionalism is correct, and cognitive systems can stretch out into the world, then living systems cannot also be co-extensive with cognitive systems, as per enactivism. Here is Clark and Kiverstein (2009) diagnosing the situation:

If living systems and cognitive systems are identical, both systems must have boundaries that coincide. However, the boundaries of the living systems are the physical boundaries of the organism. If extended functionalism is correct, the boundaries of cognitive systems can criss-cross the physical boundaries of the organism. This is precisely what the enactivists cannot allow. (p.2).

Extended cognition requires that cognitive systems recruit resources outside the boundary of the individual. Enactivism, however, denies this possibility. It therefore undercuts the identification of cognitive systems with extended systems by maintaining a sharp distinction between the physical boundaries of the organism and the environment, assuming also that each view is taken to be a global thesis about cognition.

Tensions (i) and (ii) emerge as a function of the opposing elements within 4E cognition. Tension (i) emerges as a result of body-centrism's emphasis on the unique contribution of the bodily structures, while tension (ii) follows from the location-neutrality of extended functionalism. The tensions are important for at least two reasons. One is that they stand to undermine the collective thrust of 4E cognition by showcasing fractures within the larger framework (Clark, 2008a). Another is that they reveal a lack of 'deep theoretical core' within 4E cognition. They expose a conceptual gap at the centre of an otherwise vibrant and animated collection of research (Clark & Kiverstein, 2009).

One constructive proposal that has been offered to resolve the first of the two tensions is Clark (2008a). Clark's suggestion is that the body plays an enabling computational role within cognitive processes that selectively impacts both conscious and non-conscious computational strategies. Clark's view is that the first tension can be resolved by

viewing bodily structures in terms of enabling different kinds of information processing. The cognitive significance of the body resides in the functional role it occupies within ‘intelligent’ organization – this is what explains the intuition that the body makes a special contribution to cognition.

Clark’s proposal, although not explicitly, also provides a solution to the ‘deep theoretical core’ problem. This is because it articulates, at least in principle, a common ‘computational/functional core’ for 4E cognition. Embedded, embodied, and extended approaches are unified by a shared emphasis on distributed functional complexes supporting cognitive activities. What is important is that bodily or environmental structures are situated within a larger computational/functional framework during investigation. Some systems will be individual bound (as per enactive and embodied cognition), while others will spread out across brain, body and world (as per extended cognition).

One problem with Clark’s response, despite its advantages, is that it fails to specify the relationship between physical mechanisms and computational systems finely enough. It fails to cash out what it is that allows the body to play its ‘enabling role’ in cognitive activities in the first place. If the body is merely one element within a larger brain-body-world complex, why should it have such a constraining and enabling role? The problem is not that Clark is wrong in proposing that the body has an enabling computational role, but that the suggestion alone does not suffice to specify what the role amounts to and why it should prove important.

Two Tensions Resolved

In what follows, we argue that Clark’s proposal can be supplemented and further developed by appealing to the notion of ‘wide computationalism’. We begin by outlining and motivating wide computationalism and then turn to showing how the view addresses each of the two tensions.

Wide computationalism is the view that some of the units of computational cognitive systems reside outside the individual (Wilson, 1994, 1995, 2004; Hutchins, 1995; Clark & Wilson, 2009; Kersten, 2016; Kersten & Wilson, 2016). Wide computationalism stakes a claim on the scope of physical systems, processes, and components that are capable of supporting computational analysis. A wide computational perspective opens up the possibility of exploring computational units that include the brain and aspects of the beyond-the-head environment.

Wide computationalism gains a theoretical foothold via the location neutrality of computational individuation. Since formal systems are indifferent to physical medium and computation is a formal system, it is possible that at least some states and processes relevant to a computational system may reside outside the individual. Nothing in the method of computational individuation precludes the possibility of wide computational systems.

Traditionally, wide computationalism has been committed to what some call “causal mapping accounts” of

computation (Chalmers, 1994; Chrisley, 1995). Causal mapping accounts maintain that in order for a physical system to implement an actual computation there must be a mapping of computational states to physical states such that transitions between the physical states result in corresponding transitions between the computational states. Causal mapping accounts, whether wide or narrow, articulate the conditions for ascription of computational implementation in terms of isomorphic mappings between computational descriptions and physical descriptions via transitions between physical states.

More recently, some have argued that wide computationalism should adopt a ‘mechanistic’ approach to computation (Kersten, 2016). Wide mechanistic computation differs from causal mapping formulations in that it frames the conditions of concrete computation in terms of functional mechanisms (Milkowski, 2013, 2015; Piccinini, 2015; Dewhurst, 2016). Mechanistic accounts maintain that concrete computations occur wherever there is a physical system that has an organization of spatiotemporal components such that it computes an abstract function in virtue of manipulating medium-independent vehicles. The mechanistic approach emphasizes functionally integrated systems that compute at least one abstract function via vehicle manipulation.

The wide account of computation extends the mechanistic reasoning to brain-body-world systems. It maintains that whether or not functional mechanisms, ones that process medium-independent vehicles, are constituted by spatiotemporal components squarely localized within the individual or crisscrossing into the world is an *a posteriori* question. Since the mechanistic conditions on concrete computations are medium and location neutral, the question of wide computational systems is an open one – some physical computing cognitive systems may be ensconced within the body, while others may be spread out over brain, body and world.

In addition to its theoretical plausibility, wide computationalism also gains support from studies in animal and human psychology. There is a natural set of phenomena productively studied by wide computationalism.

Research in form perception, for example, shows that formal primitives in the environment are often relevant to computationally explaining the construction of complex, internal representations. Wilson (1994, 1995) takes such work to be indicative of a wide computational system, as it acknowledges the unique computational role of states beyond the individual within perceptual processing.

Research on the spatial navigation of bats has also been used to support wide computationalism. Kersten (2016), for instance, argues that bats’ navigation system instantiates a wide computational system in virtue of employing a functional mechanism that spans the brain, body and world, and which processes medium-independent acoustic vehicles. Bat morphology, acoustic signals and neural processing conspire to support object detection along vertical planes using a wide computational system (MacIver, 2009).

Finally, Hutchins (1995) has argued for the presence of wide computation in the context of ship navigation. Hutchins' claim is that members of a navigation team carry out computational tasks that extend beyond the local actions of individual team members, which is indicative of a wide computational system.

The central message is that wide computational systems are not only theoretically plausible, but they are actually implemented in a number of cases. Research in human and animal psychology delivers several examples of concrete computational systems that extend beyond the boundary of the individual. In what follows, we adopt the wide mechanistic account of computation, though, for ease of exposition, we refer to it simply wide computationalism.

Consider how wide computationalism might address the first of the two tensions. Noë (2004), for instance, writes: "If perception is in part constituted by out and out possession and exercise of bodily skills...then it may also depend on our possession of the sorts of bodies that can encompass those skills, for only creatures with such a body could have those skills" (p.25). Noë and other body-centrists are at pains to highlight the role and contribution of bodily actions in cognitive processes. Such considerations motivate the claim that bodily-structures are constitutive of cognitive processes.

The wide mechanistic account can accommodate these types of considerations by focusing on functional mechanisms. A system is a functional mechanism when it consists of a set of spatiotemporal components that contribute to the system's overall function in virtue of the organization and interaction of its component parts (Piccinini, 2015, p.119). For the wide computationalist, the set of bodily and neural structures responsible for delivering visual perception are the functional mechanism that carry out the larger, computational task under investigation. This means that the claims of the body-centrist can be reframed in terms of 'wide' functional mechanisms; these are mechanisms whose component parts are spread out over internal and environmental elements (see Menary, 2007, ch.2). The constitutively embodied systems, such as in the case of vision, are the wide functional mechanisms localized to the body-brain complex. One way to understand body-centric theorists, then, is as making fine-grained statements about wide functional mechanisms.

Consider the extended functionalist side of the equation. In applying the method of computational analysis to world-individual spanning systems, wide computationalism maintains a commitment to the location neutrality of cognition. What matters for the wide computationalist is the functional capacity being investigated, not the physical medium through which it is realized. The implication is that extended functionalism's emphasis on medium independence and abstract analysis is preserved within wide computationalism. The view retains the abstract form of analysis crucial to the extended functionalist.

A resolution to the first tension is in sight. By re-envisioning body-centrism in terms of the implementation of

wide computational systems space is opened up for the tight, causal integration of bodily and neuronal processes in support of cognitive processes (i.e., wide functional mechanisms) and the locational neutrality of computational individuation (i.e., wide computational analysis). The special status of bodily structures turns out to be species of a more general class of wide mechanistic systems. The only difference is that whereas some wide functional mechanisms are instantiated within individuals (as highlighted by body-centrism), others are instantiated by the brain, body and world (as highlighted by extended functionalism).

Consider the second tension. Enactivism, of the autopoietic variety, was unimpressed by the blurring of the organism/environment boundary within extended functionalism. The enactivist claimed that if cognitive systems were autopoietic systems then it followed that cognitive systems could not be extended, the underlying assumption being that autopoietic systems were organism bound.

One route to reconciling the two views is to show that autopoietic theory is compatible with computationalism more generally, assuming also that extended functionalism is congenial to wide computationalism more generally (see Wilson, 2004). Insofar as wide computationalism is a species of computationalism, the compatibility of computationalism with autopoietic systems theory suffices to show the compatibility of extended functionalism with autopoietic theory.

One reason to think that autopoietic theory is compatible with computationalism is a common emphasis on mechanistic explanation (Machamer, Darden, & Craver, 2000; Bechtel, 2008). Consider, for instance, what Maturana and Varela (1980) write about the methodology of autopoietic theory:

An explanation [of autopoietic theory] is always a reformulation of a phenomenon showing how its components generate it through their interactions and relations...the elements used in the explanations are bodies and their properties...they are relations and their relations, independently of the nature of the bodies that satisfy them...This mode of thinking is not new, and is explicitly related to the very name of mechanisms. (pp. 75-76).

Autopoietic theory is, at root, a functional or mechanistic approach to explanation. What matters is that systems are explained in terms of the interaction of component parts with each other and the environment, regardless of whether the systems under investigation are biological or cognitive in character. There is a functional/mechanistic mode of explanation underlying autopoietic theory.

Compare this with computationalism. Computational analysis involves detailing how the arrangement and interaction of various components conspire to process information bearing vehicles. The mechanistic account requires showing how concrete, functional mechanisms compute medium-independent vehicles in virtue of

processing some portion of their physical structure. Computational analysis is also a species of mechanistic explanation (Milkowski, 2013).

There turns out, then, to be little incompatibility between computationalism and autopoietic theory. Both approaches employ something like a mechanistic explanatory strategy when investigating phenomena, although autopoietic theory is generally pitched at lower-level biological phenomena than computationalism. The common focus on mechanistic explanation ensures that there will be an overlapping set of phenomena productively studied by both approaches.

Why, then, the perceived tension? One reason is that many of the systems analyzed by autopoietic theory are located or ensconced within the organism. Most autopoietic systems are contained within the organism as a matter of empirical fact. However, this alone does not imply that all systems must be analyzed so as to localize within the individual; parts of the environment may still come to be included within the larger analysis. In principle, autopoietic theory, similar to computationalism, is location neutral.

Another reason for the perceived tension is that some within the enactivism literature assume that computationalism implies a commitment to representation and/or information processing theories (Di Paolo, 2009). These authors assume that because computationalism entails a commitment to representation and information processing theories, and autonomous, self-regulating systems stand in contrast to these views, enactivism must be opposed to computationalism.

However, as the previous discussion of mechanistic computation illustrates, there are a number of viable accounts of computation that are minimal in their commitment to representation or information processing theories (Stich, 1983; Egan, 1995; Piccinini, 2008). There is little reason to think that an opposition to representation and information processing commits enactivism to an opposition to computationalism more generally.

The point to note is that because wide computationalism places a greater emphasis on the way in which computational processes are grounded in particular physical mechanisms it creates a link between the body-centric claims of enactivism and embodied cognition, on the one hand, and the functional considerations of extended cognition, on the other. It is in virtue of analyzing mechanisms from several vantages – some quite fine-grained, others quite coarse-grained – that wide computationalism is able to mediate the competing claims of the three strands. The view offers a philosopher's stone of sorts through which to translate the various claims of the enactivist, embodied and extended theorist.

To be a bit more specific, tension (i) is resolved by the fact that wide computationalism allows the claims of the body-centrist to be reframed in terms of wide functional mechanisms, while tension (ii) is resolved by the fact that wide computationalism, in virtue of being a species of computationalism more generally, is theoretically compatible with the basic methodology of autopoietic theory. The success of wide computationalism, therefore,

stems from the fact that it retains several of the central insights and elements that prove important to embodied, enactive and extended theorists.

One interesting implication of the preceding analysis is that it reveals what might be called the 'computational/mechanistic' core of 4E cognition. The discussion of wide computationalism goes some way to showing that something akin to 'computational/mechanistic' explanation may underwrite a fair amount of 4E cognition. Mechanistic explanations, which are explanations of systems in terms of the activities and organization of component parts, turn out to be important not only for cognitive science and psychology more generally, but for 4E cognition specifically (Craver, 2006). This point is only provisional, of course. But it does point to a promising future line of inquiry. A continued focus on the 'computational/mechanistic' underpinnings of 4E cognition may well serve to further clarify and unify the field as a whole.

Conclusion

The preceding discussion is only the first step in a larger analysis. More still needs to be said. Nonetheless, the discussion is important because it offers one route to resolving the two tensions troubling 4E cognition. What's more, the discussion shows that a renewed focus on computationalism, particularly of the wide variety, may have key role to play in illuminating the conceptual foundations of 4E cognition. This result is both interesting and novel, as discussions of 4E cognition sometimes eschew mention of computation. By showing that a form of computationalism provides a theoretically flexible yet robust vehicle through which to understand and translate the various strands of 4E cognition, the current discussion provides not only a partial vindication of wide computationalism, but also helps to shed light on an important set of issues facing a growing research programme within cognitive science.

References

- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Malden: Blackwell.
- Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. London, Routledge.
- Chalmers, D. (1994). On implementing a computation. *Minds and Machines*, 4(4), 391–402.
- Clark, A. (2005). Intrinsic content, active memory, and the extended mind. *Analysis*, 65, 1–11.
- Clark, A. (2008a). Pressing the Flesh: A Tension in the Study of the Embodied, Embedded Mind? *Philosophy and Phenomenological Research*, 76(1), 37–59.
- Clark, A. (2008b). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York: Oxford.
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58, 7–19.

- Clark, A., & Kiverstein, J. (2009). Introduction: Mind Embodied, Embedded, Enacted: One Church or Many? *Topoi*, 28, 1-7.
- Craver, C. (2006). When Mechanistic Models Explain. *Synthese*, 153(3), 355-76.
- Chrisley, R.L. (1995). Why everything doesn't realize every computation. *Minds and Machines* 4: 403-430.
- Dewhurst, J. (2016). Review of Physical Computation. *Philosophical Psychology*. DOI: 562 10.1080/09515089.2016.1150450.
- Di Paolo, E. A. (2009). Extended Life. *Topoi*, 28, 9-21.
- Egan, F. (1995). Computation and Content. *Philosophical Review*, 104, 181-203.
- Gallagher, S. (2005). *How the body shapes the mind*. Oxford: OUP.
- Haugeland, J. (1998). Mind Embodied and Embedded. In J. Haugeland (ed.), *Having Thought: Essays in the Metaphysics of Mind* (pp. 233- 267). Cambridge, MA: Harvard University Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Hutto, D. and Myin, E. (2013). *Radicalizing Enactivism: Basic minds without content*. Cambridge, MA: MIT Press.
- Lakoff, G. (1980). The Metaphorical Structure of the Human Conceptual System. *Cognitive Science*, 4, 195-208.
- Machamer, P. K., Darden, L., & Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science*, 67: 1-25.
- MacIver, M.A. 2009. Neuroethology: From morphological computation to planning. In (eds.) P. Robbins and M. Aydede, *The Cambridge Handbook of Situated Cognition* (pp.480-504). New York, Cambridge University Press.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: the realization of the living*. Dordrecht, Holland: Kluwer Academic Publishers.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave Macmillan.
- Menary, R. (2010). Introduction to the special issue on 4E cognition. *Phenomenology and Cognitive Science*, 9:459-463.
- Machamer, P. K., Darden, L., & Craver, C. (2000). Thinking about Mechanisms. *Philosophy of Science*, 67, 1-25.
- MacIver, M.A. (2009). Neuroethology: From morphological computation to planning. In (eds.) P. Robbins and M. Aydede, *The Cambridge Handbook of Situated Cognition* (pp.480-504). New York, Cambridge University Press.
- Milkowski, M. (2013). *Explaining the Computational Mind*. Cambridge, MA.: MIT Press.
- Milkowski, M. (2015). Computational mechanism and models of cognition. *Philosophia Scientiae* 18(3): 1-14.
- Noë, A. (2004). *Action in Perception*. MIT Press.
- Kersten, L., & Wilson, R.A (2016). The Sound of Music: Externalist Style. *American Philosophical Quarterly*, 53(2), 139-154.
- Kersten, L. (2016). A Mechanistic Account of Wide Computationalism. *Review of Psychology and Philosophy*, (online), 1-17.
- Kirsch, D. (2009). Problem solving and situated cognition. In (eds.) P. Robbins and M. Aydede, *The Cambridge Handbook of Situated Cognition*, (pp. 264-306). New York, NY: Cambridge University Press.
- Kirsh, D., & Maglio, P. (1995). On distinguishing epistemic from pragmatic actions. *Cognitive Science*, 18, 513-549.
- O'Regan, J. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 25(4), 883-975.
- Piccinini, G. (2007). Computing mechanisms. *Philosophy of Science*, 74(4): 501-526.
- Piccinini, G. (2008). Computation without Representation. *Philosophical Studies*, 137(20), 205-241.
- Piccinini, G. (2015). Physical Computation: A Mechanistic Account. Published to Oxford Scholarship Online. doi: 10.1093/acprof:oso/9780199658855.001.0001.
- Stich, S. (1983). *From folk psychology to cognitive science*. Cambridge: MIT Press.
- Thompson, E. (2007). *Mind in Life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.
- Varela, F., E. Thompson, and E. Rosch. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge: MIT Press.
- Weber, A., & Varela, F. J. (2002). Life after Kant: Natural purposes and the autopoietic foundation of individuality. *Phenomenology and the Cognitive Sciences*, 1(2), 97-125.
- Wheeler, M. (2009). *Reconstructing the Cognitive World*. MIT Press.
- Wheeler, M. (2010). In defense of extended functionalism. In R. Menary (Ed.), *The Extended Mind* (pp. 245-270). MIT Press.
- Wilson, R. (1994). Wide computationalism. *Mind*, 103(4), 351-372.
- Wilson, R. (1995). *Cartesian psychology and physical minds: Individualism and the sciences of the minds*. Cambridge: Cambridge University Press.
- Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences*. Cambridge: Cambridge University Press.
- Wilson, R. (2010). Extended vision. In (eds.) N. Gangopadhyay, M. Madary, and F. Spicer, *Perception, action and consciousness* (pp.277-290). New York: Oxford University Press.
- Wilson, R.A., and Clark, A. (2009). How to situate cognition: Letting nature take its course. In (eds.) M. Aydede and P. Robbins, *The Cambridge handbook of situated cognition* (pp.55-77). New York: Cambridge University Press.
- Wilson, R. A., & Foglia, L. (2016). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.