

Introduction to “The Material Bases of Cognition”

Kenneth Aizawa

Received: 12 May 2013 / Accepted: 13 May 2013 / Published online: 18 May 2013
© Springer Science+Business Media Dordrecht 2013

Twenty years ago, a special edition of a journal devoted to the material bases of cognition would likely have focused on the divide between type and token identity theories of cognition, reductionism (of some sort) versus functionalism (of some sort), or neuroscientifically-inspired versus computationally-inspired approaches to cognition. Those issues remain, but the rise of mechanistic explanation, the introduction of dynamical systems approaches to cognition, and the prominence of claims that the mind is embodied or even extended has changed the landscape. For example, where 20 years ago there was a consensus that the material bases of cognition in some sense reside in the brain, this consensus has since broken down. Some philosophers and cognitive scientists now vigorously defend the bold claim that cognitive processes have their material basis in the non-neural body, and sometimes even the extracorporeal environment, as well as the brain.¹

The papers in this special issue track this developmental arc. The collection begins with papers by Theurer and Gillett on the metaphysics of material bases. Theurer’s paper, “Compositional explanatory relations and mechanistic reduction,” focuses on a putative feature of reduction (or reductive explanation), namely, whether the relation, “X is a material basis for Y” (and hence, “X mechanistically explains Y”) is a transitive relation. And, if so, how is this possible? Gillett’s paper, “Constitution, and multiple constitution, in the sciences,” offers a theory of what it is for, say, a neuron in a brain, but not a bullet lodged in the brain, to be a part of the material basis of the brain—and, therefore, an account of what it might be for, say, hands or tools to be parts of the material basis of a cognitive system.

¹ See, for example, Beer (2003), Clark and Chalmers (1998), Rowlands (1999), and Haugeland (1998).

K. Aizawa (✉)
Department of Philosophy, Rutgers, The State University of New Jersey, Newark Campus,
Conklin Hall, 175 University Avenue, Newark, NJ 07102, USA
e-mail: ken.aizawa@gmail.com

The papers by Adams and Garrison, and Shapiro, concern themselves with the nature of cognition. In “The mark of the cognitive,” Adams and Garrison ask what distinguishes cognitive processes from non-cognitive processes?² They also explain how a familiar theory of the cognitive meets some of the challenges of embodied and extended cognition. In “Dynamics and Cognition,” Shapiro explores some anti-cognitivist, anti-representationalist arguments that have appeared in dynamical systems approaches to cognitive science.

Huneman and Haselager, for their parts, explore new directions in the debates over the material bases of cognition. Huneman’s paper, “Causal parity and externalisms,” provides a critical treatment of the apparent commonalities between (a) the idea that cognitive science should study cognitive systems that are extended beyond the boundaries of the organism and (b) the idea that evolutionary and developmental biology should study developmental systems that are extended beyond the boundaries of organisms.³ In “Did I do that?,” Haselager makes two principal moves. First, where much of the discussion in the embodiment literature has focused on cognition, Haselager focuses on one’s sense of agency, i.e., the sense that one is, or is not, the author of some action.⁴ Second, Haselager offers a suggestive account of how certain tools, namely, wheelchairs controlled through the assistance of artificial intelligence and brain–computer interfaces (BCI), causally influence one’s sense of agency. What Haselager implicitly proposes is that there can be scientifically and philosophically interesting work showing how one’s tools can causally influence one’s sense of agency, so that one does not need to go so far as to claim, for example, that one’s use of tools constitutes one’s sense of agency.

The remainder of this introduction will provide more commentary on the papers. It will aim to provide an overview of their principal themes, ideas, and arguments, but it will also try to amplify or clarify certain conclusions and draw attention to further inter-relations between the papers. Such an essay might be more valuable than abstracts of the individual papers. These comments are, of course, the editor’s alone and should not be construed as representative of any of the individual authors’ views.

Returning to Theurer’s (2013) paper, we may note that it centers on a scientifically possible scenario. One might find, schematically speaking, cases in which facts at level 2 are apparently mechanistically explained by facts at level 1 and in which facts at level 1 are apparently mechanistically explained by facts at level 0. In such cases, do the facts at level 0 explain the facts at level 2?⁵ Perhaps

² Much has been written on this topic, including Adams and Aizawa (2001, 2008), Rowlands (2009), and Rupert (2009).

³ As a philosophical introduction to developmental systems theory, one might begin with Griffiths and Gray (1994). For comparisons of dynamical systems theory with extended mind, see Schulz (2013) and Wilson (2004, 2005).

⁴ See Colombetti (2007), Krueger (2009), and Krueger and Overgaard (2012), for a defense of embodied emotions, i.e., the view that emotions are constituted by, and not merely expressed by, bodily actions, such as the clenching of one’s fists or the grinding of one’s teeth.

⁵ Many of those party to this debate have taken to using “reduction” to describe this kind of “interlevel” explanation, so that one can write almost interchangeably about “reduction,” “mechanistic explanation,” “mechanistic reduction,” and “interlevel explanation.” For simplicity, I will stick with (interlevel) explanation.

vision science gives us an actual case. Some age-related changes in color discrimination are apparently explained by the yellowing of the lens and the yellowing of the lens is apparently explained by changes in the biochemical composition of the lens. So, do the changes in the biochemical composition of the lens explain the changes in color discrimination? What is to be said about such scenarios? Bickle (2003, 2006) maintains that level 0 can (directly) explain level 2 without appeal to level 1. So, Bickle thinks that one can (directly) explain the age-related changes in color discrimination by appeal to the biochemical changes in the lens. Theurer endorses this claim, but offers new arguments to support it. By contrast, Bechtel (2009) maintains that explanations are always just one level down, so that even when level 0 explains level 1 and level 1 explains level 2, level 0 does not explain level 2.⁶ So, Bechtel thinks that level 1 is necessary for an explanation of level 2. Thus, according to Bechtel, one cannot (directly) explain the age-related changes in color discriminations by appeal to the biochemical changes in the lens.

To support her view, Theurer proposes that we embrace an idea about explanation that has largely been lost, namely, that certain types of explanation are transitive. If interlevel mechanistic explanations are based upon ontological relations among things in the world, then if the relevant ontological relation forming the backbone, so to speak, of an interlevel mechanistic explanation is transitive, then we would find that if level 0 explains level 1 and level 1 explains level 2, then level 0 also explains level 2. What, then, might such a transitive ontological backbone for interlevel explanatory relations be?

Theurer considers several options. Identity between higher and lower level entities was once a popular proposal, but it has features that make it ill-suited to be this explanatory backbone. For one thing, properties at level 0 explain a property at level 1, but not vice versa. So, mechanistic explanation is an asymmetric relation, where identity is a symmetric relation. For another thing, since many properties at level 0 explain a property at level 1, mechanistic explanation is a many-to-one relation, where identity is a one-to-one relation. Classical mereological composition, as found for example in Lewis (1991), will also not do. For one thing, mereological composition allows for every object to be a part of itself, but interlevel explanations do not allow for such reflexivity. Instead of identity and mereological compositionality, Theurer finally appeals to the theory of material compositionality developed in Koslicki (2008). This is a theory that has many of the features that are needed for interlevel explanation, including transitivity. Theurer's endpoint, thus, complements Gillett's paper on the compositional relations between individuals in the sciences.

Gillett's (2013) contribution to this volume is an installment in his larger project of providing an account of the types of compositional relations underlying scientific explanations. (See, for example, Aizawa and Gillett (2009).) Gillett assumes interlevel mechanistic explanation is based on compositional relations and seeks to articulate the nature of these concepts. At the risk of some anachronism, we might say that, since the Scientific Revolution, (a) individuals have been explained using smaller individuals taken to be their parts, (b) the properties of such individuals

⁶ See Bechtel (2009).

have been explained using the properties of their parts, and (c) the processes grounded by these individuals have been explained using the processes grounded by their parts.⁷ Twentieth Century science, especially biology, has vindicated these ideas in dramatic fashion. As a clear example, consider the eukaryotic cell. On Gillett's theory, (a) the cell is constituted by molecules (among other entities), (b) the cell's property of rigidity is realized, in part, by the properties of its cytoskeletal molecules, and (c) a cell's process of moving is implemented, in part, by the processes grounded by the cytoskeletal molecules. This integration of individuals, properties, and processes, and the compositional relations between these entities at different levels, is central to Gillett's account of compositional relations in the sciences.

Omitting many ideas and details, Gillett's paper weaves properties and processes into a theory of the part-whole (or constitution) relations of individuals in the sciences, providing an account of what it is for an individual to constitute a "working part" of another individual. For instance, under what conditions are mitochondria working parts of a eukaryotic cell? On Gillett's account, we must, first, ask whether mitochondria are spatially contained within the cell. Second, we must ask whether the properties and relations of the mitochondria realize, in part, the properties and relations of the cell (but not vice versa). And also, consequently, whether the processes grounded by mitochondria implement the processes grounded by the cell (but not vice versa). Insofar as the answers to these questions are "yes," then mitochondria are parts of a cell. (Is a bullet lodged in the brain a "working part" of the brain? If the answer to any one of the corresponding questions is "no," then the bullet is not.)

Among its many merits, Gillett's work on scientific compositional relations deserves greater attention in the extended and embodied cognition debate. To understand why this is, we need to note two features of the debate. First of all, advocates of extended and embodied cognition have often claimed that cognitive processes do not merely causally interact with non-cortical bodily and environmental processes; instead, they are sometimes constituted by non-cortical bodily and environmental processes. The distinction between these two views is perhaps best known through the labels of (HEC) and (HEMC), introduced in Rupert (2004):

(HEC) human cognitive processing literally extends into the environment surrounding the organism, and human cognitive states literally comprise-as wholes do their proper parts-elements in that environment; in consequence, while the skin and scalp may encase the human organism, they do not delimit the thinking subject (Rupert 2004, p. 389).

(HEMC) cognitive processes depend very heavily, in hitherto unexpected ways, on organismically external props and devices and on the structure of the external environment in which cognition takes place (Rupert 2004, p. 393).

A second feature of the debate has been the common (though not universal) presupposition that if X is connected to a cognitive system Y by means of a causal

⁷ Powers could also be added to the mix, but these are omitted here for simplicity.

connection of *the right kind*, then X is a part of that cognitive system Y.⁸ In short, some sort of causal relation is frequently thought to provide a sound basis for inferring a constitutive relation.

Gillett has important contributions to make to both of these issues. First, advocates of extended and embodied cognition have yet to articulate a theory of the putative difference between causation and constitution—the distinction that separates (HEC) and (HEMC)—that is central to their approach.⁹ Gillett, however, provides an account of the relation of constitution and its differences from causal (or productive) relations.¹⁰ On the second issue, Gillett provides an alternative account of the basis of constitutive relations that illuminates their connections to causal (or productive) relations. On Gillett’s approach, X does not become a part of an individual Y in virtue of causal (productive) relations that X itself bears to Y.¹¹ Instead, X is a part of Y in virtue of entering into what we might summarize as X’s co-productive, co-realizing and co-implementing relations with Y’s other parts. Gillett’s proposals, thus, bridge the mechanistic explanation literature and the embodied, extended cognition literature.

Both the Adams and Garrison (2013) paper and the Shapiro (2013) paper probe what advocates of extended cognition think of (a) the relationship between cognition and behavior, and (b) the role of representation in cognition. Adams and Garrison begin by drawing attention to debates (other than extended cognition) for which a theory of the cognitive would be valuable, namely, discriminating between beings (e.g., machines and organisms) with cognitive lives and beings without them. Adams and Garrison next argue that behavior (mere bodily movement) does not wear cognition on its sleeve. One can kick out one’s leg as a non-cognitive patellar reflex or one can kick out one’s leg intentionally. What is the difference? Adams and Garrison argue that the latter sort of case is a matter of acting for one’s own reasons.¹² They also argue that Rowlands (2009) does not provide an adequate account of the cognitive. Rowlands maintains that to be a cognitive process involves something being a representation for a subject. Adams and Garrison, however, argue that being a subject apparently presupposes having cognitive processes. So, rather than providing an account of cognitive processes, Rowlands assumes one. Finally, Adams and Garrison propose a way to avoid the conclusion, advanced in Turner

⁸ See van Gelder (1995), Clark and Chalmers (1998), and Haugeland (1998), for variants of this idea. See Adams and Aizawa (2008), Chapter 6 and 7 for a critical discussion. Note the shift between talk of cognitive processes and cognitive systems. Adams and Aizawa (2008) treat the hypothesis that cognitive processes are extended and the hypothesis that cognitive systems are extended as substantially distinct claims (the former being stronger than the latter), where Rupert (2009) and Wilson and Clark (2009), for instance, treat them as largely readily inter-definable variants.

⁹ Indeed, some advocates of extended and embodied cognition have even resisted the distinction upon which the debate appears to rest. See Hurley (2010) and Rockwell (2010).

¹⁰ See Gillett (2007) and Craver (2007).

¹¹ Gillett, in fact, presents objections to the version of this idea found in Mellor (2008).

¹² This account of the cognitive potentially diverges from that given in Adams and Aizawa (2001, 2008). These earlier collaborative works focused on the cognitivist idea that cognition is a matter of specific sorts of manipulations of non-derived representations. For Adams and Garrison, having a reason requires having non-derived representations, but it is unclear what restrictions, if any, acting for a reason places on the character of representation handling.

(2013), that when a colony of termites repairs its mound the termites constitute a cognitive “superindividual”. Adams and Garrison claim that the termite collective does not repair for a reason, hence that it is not a cognitive “superindividual”.

Shapiro’s “Dynamics and Cognition” explores two issues concerning the relationship between the mathematics of dynamical systems theory and cognition. It, first, focuses on the claim that dynamical systems that model cognition consist of coupled equations that include variables ranging over a brain, a body, and an environment. As a possible example of this kind of dynamical system, there is the way in which Kelso (1995) seeks to understand how the direction of an agent’s finger wagging changes with increases in the frequency of the wagging.¹³ But why should we think finger wagging is cognition? This looks like what would long have been taken to be behavior. The answer, for some apparently, is that we should use “cognition” as a term for behavior. As Chemero puts it, “cognitive scientists ought to try to understand cognition as intelligent behavior” (Chemero 2009, p. 25). If this terminological shift is the new dynamicist’s game, then there would not be much empirical interest in the playing.

The second issue in Shapiro’s paper concerns the role of representations in dynamical systems theory. In truth, there are many forms of anti-representationism and many anti-representationalist arguments to be found in cognitive science these days, so that the discussion must be limited to two. First, Beer (2003) proposes that representations must have discrete, linguaform vehicles, but that there are no such things in the models he has developed. Yet, one might reply that representations need not have this character. The firing of individual neurons might serve as non-linguaform representations of features of the environment, such as luminance. In addition, Ward and Ward (2009) showed how a connectionist network of the sort that Beer had studied can develop neuron-like structures that correspond an object’s shape. These, too, would seem to be representations without discrete, linguaform vehicles. Second, sometimes it is maintained that representations are dispensable when an agent is in constant contact with the environment it cognitively copes with. “The world is its own best model” is an oft cited shorthand for this idea.¹⁴ This objection to representation, however, appears to be of limited scope. In many cases, agents are not in constant contact with the things they are thinking about. As a simple and vivid case, Shapiro, sitting in his office in Wisconsin, can evidently visually imagine walking through his childhood home in New Jersey without any current contact with his childhood home.

Importantly, Shapiro, like Adams and Garrison, draws attention to the familiar idea that cognitive processes are among the causes of behavior. This familiar idea suggests a useful alternative to asking for a full-blown theory of the mark of the cognitive.¹⁵ Perhaps we can ask of the advocates of embodied or extended cognition whether they maintain that one’s cognitive processes are non-behavioral causes of

¹³ See Kelso (1995).

¹⁴ See Brooks (1997).

¹⁵ Clark (2010) and Wilson (2010) dismiss the putative need for such a theory. Clark proposes that there may not be anything that marks off the cognitive from the non-cognitive, where Wilson proposes that the question invites a methodologically dubious exercise in conceptual analysis. Chemero (2009, p. 212, fn. 8) also strikes a note of resistance to providing a “definition” of the cognitive.

one’s own behavior. One’s behavior at one time can, of course, cause changes in one’s own behavior, or the behavior of others, at later times, but can one’s own cognitive processes be non-behavioral influences on one’s own behavior? If one shies away from the question about what distinguishes cognitive processes from non-cognitive processes, then perhaps one can at least answer this alternative question. Perhaps this is a useful way to advance the extended cognition debate without worrying about such meta-theoretic matters as whether having a theory of cognition is a matter of giving definitions or conceptual analysis.

As noted above, Huneman’s “Causal parities and externalisms,” examines the apparent similarities between the extended mind hypothesis in cognitive science and developmental systems theory in biology. Although a diversity of arguments might be offered in support of these perspectives, Huneman focuses (roughly) on the way in which consideration of causal relations between the mind and genes, on the one hand, and the body and environment, on the other, make a case for the existence of a kind of unitary organism-environment system. All parties to the recent debate over extended cognition in cognitive science agree that cognitive processes exist in a complicated causal nexus spanning brain, body, and world, and all parties to the recent debate over developmental systems theory in evolutionary biology agree that genes exist in a complicated causal nexus spanning cell nuclei, body, and world. Given the existence of such a causal nexus, however, there are two types of questions one might ask: (1) What are the types of the causes in the nexus? and (2) What are the relative “weights” of the causes in the nexus? These are the issues at the heart of the debates over extended cognition and developmental systems theory.

To a first approximation, the debate over extended cognition has centered on the first question, regarding the types of the causes in the nexus. As Huneman puts it, “The internal stuff and the external stuff might be equivalent *qua* causal, but one might also want to know whether the internal and the external constitute the same kinds of causes.” In particular, are they cognitive causes? More concretely, the debate has been about whether something like the whole process of writing a memo in a notebook is (under certain conditions) the same type of cause of future behavior as is the process of committing something to brainy memory. Are both processes cognitive processes?¹⁶ Supporters of extended cognition believe that the causal processes extending out to and from the notebook are (under certain conditions) cognitive causal processes, where opponents of extended cognition believe that these are typically non-cognitive causal processes. Here again, it would seem we need a “mark of the cognitive” to adjudicate the matter.

In contrast to the extended cognition debate, much of the debate in developmental systems theory has focused on the second question, namely, the relative importance of genetic and environmental causes in organismal development and evolution. Advocates of the developmental systems theory maintain that much work on the evolutionary synthesis has assigned an unwarranted primacy to the role of genes in the lives of organisms. This means that there is a disanalogy between

¹⁶ This way of putting the issue presupposes a version of extended cognition according to which both intracranial and transcranial processes are cognitive. Clark and Chalmers (1998) provides an instance of this type of view. Other versions of extended cognition deny that there is anything intracranial that might be called cognition. Haugeland (1998) and Chemero (2009) provide instances of this type of view.

extended cognition and extended life. To change terminology somewhat, so-called “brain-centrism” is the idea that certain types of processes, namely, cognitive processes, are (typically) found only in the brain. “Gene-centrism,” by contrast, is idea that the gene is of overriding importance in development and evolution.

Within this framework, Huneman explores two possible ways to try to justify brain-centrism and gene-centrism. First, the brain and the gene might bear non-derived content, that is, they might bear content that does not depend on something like social conventions.¹⁷ Second, brain-centrism and gene-centrism do not describe how brains and genes necessarily work, but merely how they contingently happen to work. The brain is contingently the locus of cognitive processes and the gene is contingently the most important causal factor in evolution. Ultimately, Huneman ends up with a position somewhat sympathetic to brain-centrism, but less committed to a position on gene-centrism.

To return, finally, to Haselager’s paper, we might depart from its order of exposition. Let us begin with brief descriptions of some recent experimental work regarding ways in which a subject’s sense of agency—one’s sense of having performed some act—can be experimentally manipulated. On the one hand, there are *illusions of control*, wherein one has the false sense that one is doing something and, on the other, there are *automatisms*, wherein one is genuinely doing something without sensing that one is. While illusions of control and automatisms are infrequent natural occurrences, there are ways in which these phenomena may be experimentally induced. For example, in one of the so-called “helping hands” experiments Haselager describes, a participant watches herself in a mirror, while an accomplice hidden behind the participant, extends her arms forward on either side of the participant.¹⁸ In one condition, the participant hears an instruction (e.g., “Give the ok sign”) that the accomplice subsequently performs, whereas in another condition, the participant hears nothing before the accomplice performs the action. Participants report a stronger sense of agency in the cases where they hear the instruction the participant subsequently performs than in the cases where they do not hear the instruction. In addition, Haselager describes cases of “facilitated communication” as instances of automatisms. In a putative facilitated communication situation, a person referred to as the “facilitator” may support the hand of a communicatively impaired individual. The “facilitator” then tries to use input from the impaired individual to type messages from the individual. Yet, when the “facilitator” and the individual wear headphones, hence can be given distinct messages, the “facilitator” types what she hears, not what the impaired individual hears. Nevertheless, the “facilitators” have the sense that they are communicating the message of the impaired individual.

While cases of illusions of control and automatisms are generally limited to experimental contexts, they appear to take on a more practical cast in the light of the development of BCIs complemented by artificial intelligence devices (ID). In the relevant kind of BCI, brain activity is measured by something like

¹⁷ For a discussion of the idea that non-derived content might be an essential feature of cognitive processes, see, for example, Adams and Aizawa (2001, 2008) and Clark (2005).

¹⁸ Cf. Wegner et al. (2004).

electroencephalography (EEG) and used to control, say, the movements of a wheelchair. This technology, thereby, might enable a paralyzed individual to move about without the use of their limbs. Given the limitations on the reliability of measurements of brain activity and the relatively long times needed to collect brain activity data, however, there are currently efforts to use intermittent robotic interventions to prevent mishaps, such as having wheelchair users collide with tables, chairs, and other people. BCI-ID technology, thus, threatens to create real world cases of illusions of control and automatism. Users of this technology might have the sense that they are in control of their movements, when in fact the ID is, or they have the sense that the ID is in control, when it is not. Haselager notes that, farther into the future, BCI-ID may raise issues of legal responsibility. In the U.S. Model Penal Code, for example, a person is not guilty of certain offenses if the relevant body movement is not a product of the person’s determination. Thus, insofar as one’s sense of agency is linked to one’s determination, it appears that if one lacks the appropriate sense of agency, one may not bear legal culpability.

As a final reflection on the papers in this special issue, one can see one future direction for research on the material bases of cognition. There can be fruitful efforts to bring together the long tradition of work on type and token identity theories, reductionism, and scientific explanation—currently flourishing in the philosophy of science under the rubric of “mechanistic explanation” and represented here by Theurer and Gillett—with work on embodied and extended cognition, represented here by Adams and Garrison, Shapiro, Huneman, and Haselager.

References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*(1), 43–64.
- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Malden, MA: Blackwell Publishers.
- Adams, F., & Garrison, R. (2013). The mark of the cognitive. *Minds and Machines*. doi: [10.1007/s11023-012-9291-1](https://doi.org/10.1007/s11023-012-9291-1).
- Aizawa, K., & Gillett, C. (2009). Levels, individual variation and massive multiple realization in neurobiology. In J. Bickle (Ed.), *The Oxford handbook of philosophy and neuroscience* (pp. 539–581). Oxford: Oxford University Press.
- Bechtel, W. (2009). Molecules, systems, and behavior: Another view of memory consolidation. In J. Bickle (Ed.), *The Oxford handbook of philosophy and neuroscience* (pp. 13–40). Oxford: Oxford University Press.
- Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *11*(4), 209–243.
- Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive account*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Bickle, J. (2006). Reducing mind to molecular pathways: Explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese*, *151*, 411–434.
- Brooks, R. A. (1997). Intelligence without representation. In J. Haugeland (Ed.), *Mind design II* (pp. 395–420). Cambridge, MA: MIT Press.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Clark, A. (2005). Intrinsic content, active memory and the extended mind. *Analysis*, *65*(285), 1–11.
- Clark, A. (2010). Coupling, constitution and the cognitive kind: A reply to Adams and Aizawa. In R. Menary (Ed.), *The extended mind* (pp. 81–99). Cambridge, MA: MIT Press.

- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Colombetti, G. (2007). Enactive appraisal. *Phenomenology and the Cognitive Sciences*, 6(4), 527–546.
- Craver, C. (2007). *Explaining the brain*. Oxford: Oxford University Press.
- Gillett, C. (2007). Understanding the new reductionism: The metaphysics of science and compositional reduction. *The Journal of Philosophy*, 104(4), 193.
- Gillett, C. (2013). Constitution, and multiple constitution, in the sciences: Using the neuron to construct a starting framework. *Minds and Machines*. doi:10.1007/s11023-013-9311-9.
- Griffiths, P. E., & Gray, R. D. (1994). Developmental systems and evolutionary explanation. *The Journal of Philosophy*, 91(6), 277–304.
- Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Ed.), *Having Thought* (pp. 207–237). Cambridge, MA: Harvard University Press.
- Hurley, S. (2010). Varieties of externalism. In R. Menary (Ed.), *The extended mind* (pp. 101–153). Cambridge, MA: MIT Press.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- Koslicki, K. (2008). *The structure of objects*. Oxford: Oxford University Press.
- Krueger, J. (2009). Empathy and the extended mind. *Zygon*, 44(3), 675–698.
- Krueger, J., & Overgaard, S. (2012). Seeing subjectivity: Defending a perceptual account of other minds. In S. Miguens & G. Preyer (Eds.), *Consciousness and Subjectivity*. Piscataway, NJ: Ontos Verlag.
- Lewis, D. (1991). *Parts of classes*. Cambridge, MA: Basil Blackwell.
- Mellor, D. H. (2008). Micro-composition. *Royal Institute of Philosophy Supplements*, 62, 65–80. doi: 10.1017/S1358246108000581.
- Rockwell, T. (2010). Extended cognition and intrinsic properties. *Philosophical Psychology*, 23(6), 741–757.
- Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. New York, NY: Cambridge University Press.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1–19.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101(8), 389–428.
- Rupert, R. (2009). *Cognitive systems and the extended mind*. New York: Oxford University Press.
- Schulz, A. W. (2013). Overextension: The extended mind and arguments from evolutionary biology. *European Journal for Philosophy of Science*. doi:10.1007/s13194-013-0066-1.
- Shapiro, L. (2013). Dynamics and cognition. *Minds and Machines*. doi:10.1007/s11023-012-9290-2.
- Theurer, K. (2013). Compositional explanatory relations and mechanistic reduction. *Minds and Machines*. doi:10.007/s11023-013-9306-6.
- Turner, J. S. (2013). Superorganisms and superindividuality: The emergence of individuality in a social insect assemblage. In F. Bouchard & P. Huneman (Eds.), *From groups to individuals* (pp. 219–242). Cambridge, MA: MIT Press.
- van Gelder, T. (1995). What might cognition be, if not computation? *The Journal of Philosophy*, 91, 345–381.
- Ward, R., & Ward, R. (2009). Representation in dynamical agents. *Neural Networks*, 22(3), 258–266.
- Wegner, D. M., Sparrow, B., & Winerman, L. (2004). Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology*, 86(6), 838–848.
- Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences: Cognition*. New York, NY: Cambridge University Press.
- Wilson, R. (2005). *Genes and the agents of life: The individual in the fragile sciences: Biology*. New York, NY: Cambridge University Press.
- Wilson, R. (2010). Review of Robert Rupert's *Cognitive Systems and the Extended Mind*. *Notre Dame Philosophical Reviews*. <http://ndpr.nd.edu/review.cfm?id=19128>.
- Wilson, R., & Clark, A. (2009). How to situate cognition: Letting nature take its course. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 55–77). New York: Cambridge University Press.