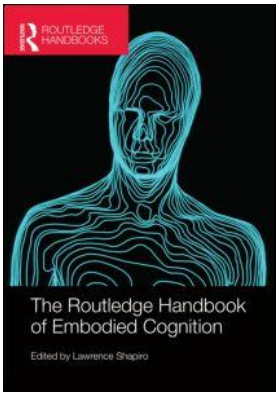

The Routledge Handbook of Embodied Cognition

Lawrence Shapiro

Extended Cognition

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9781315775845.ch3
Ken Aizawa
Published online on: 01 May 2014

PART II

Perspectives on embodied cognition

This page intentionally left blank

3

EXTENDED COGNITION

Ken Aizawa

What takes place in the environment clearly influences cognitive processing. Children raised around speakers of English typically learn to understand English, whereas children raised around speakers of German typically learn to understand German. Accountants using calculators will typically know that accounts are balanced more quickly than those who use pencil and paper. Extended cognition goes beyond such pedestrian observations. Rather than causal dependency relations between cognitive processes and environmental processes, extended cognition postulates a constitutive dependence between cognitive processes and processes in brain, body, and environment. Cognitive processes are realized, not just in the brain, but also in the body and world.

This brief chapter will focus on two types of arguments for extended cognition inspired by Clark and Chalmers (1998). First, there has been the thought that cognition extends when processes in the brain, body, and world are suitably similar to processes taking place in the brain. We might describe these as cognitive equivalence arguments for extended cognition. Second, there has been the thought that, when there is the right kind of causal connection between a cognitive process and bodily and environmental processes, cognitive processes come to be realized by processes in the brain, body, and world. We might describe these as coupling arguments for extended cognition. What critics have found problematic are the kinds of similarity relations that have been taken to be applicable or suitable for concluding that there is extended cognition and the conditions that have been offered as providing the right kind of causal connection.

Cognitive equivalence arguments

The best-known cognitive equivalence argument is Andy Clark and David Chalmers' "Inga-Otto" thought experiment.[1] In this experiment, Inga is a normal human being who hears about an exhibit at the Museum of Modern Art. She decides that she wishes to see it, then pauses to recollect where the museum is. When she remembers, she heads out to 53rd Street. Next consider Otto, an individual suffering from the early stages of Alzheimer's disease. In order to cope with the loss of his biological memory, he turns to a notebook wherein he keeps useful information, such as the location of the museum. When he hears about the exhibit, he decides he

wishes to see it, and then flips through his notebook until he finds the address. When he does, he heads out to 53rd Street.

Clark and Chalmers claim that the information stored in their respective brainy and papery resources constitutes memory for both Inga and Otto. They claim that "in relevant respects the cases are entirely analogous: the notebook plays for Otto the same role that memory plays for Inga. The information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin" (Clark and Chalmers, 1998, p. 13). Further, "Certainly, insofar as beliefs and desires are characterized by their explanatory roles, Otto's and Inga's cases seem to be on a par: the essential causal dynamics of the two cases mirror each other precisely" (ibid.). And, finally, "To provide substantial resistance, an opponent has to show that Otto's and Inga's cases differ in some important and relevant respect. But in what deep respect are the cases different?" (Clark and Chalmers, 1998, pp. 14–15).[2]

Substantial resistance through non-derived content?

Critics have offered a number of proposals to provide the kind of resistance that Clark and Chalmers demand. One contention has been that Inga's informational resources bear non-derived content, where Otto's informational resources bear derived content.[3] Non-derived content does not depend upon previously existing content, where derived content does. Stereotypically, non-derived content arises from social conventions. So, the meaning of a red traffic light, the meaning of a white flag, and the meaning of written words are paradigmatic instances of non-derived content.

It is sometimes claimed that cognitive states necessarily bear non-derived content, so that the fact that the letters and numerals in Otto's notebook presumably bear only derived content suffices to show that the letters and numerals are not constituents of any of Otto's cognitive states. Of course, bearing non-derived content may be important to cognition, even if it is not a strictly necessary property of a cognitive state. Perhaps cognitive states are natural kinds on the model of homeostatic property clusters.[4] On this analysis, cognitive states lack necessary or essential properties, such as bearing non-derived content, but they nevertheless tend to have a stable set of similarities, among which may be that they bear non-derived content. Thus, bearing derived content makes the inscriptions in Otto's notebook less like paradigmatic cognitive states. What this shows is that, in principle, there is more than one metatheoretical means by which to recognize the importance of non-derived content to being cognitive.

Dennett (1986), Clark (2005, 2010), Menary (2010), and Fisher (2009) argue that the derived/non-derived distinction is neither important nor relevant to cognitive status. Prior to the debate over extended cognition, Dennett (1986) gave a multifaceted challenge to the hypothesis of non-derived (or "original") content. He claimed that there is no distinction between derived and non-derived content, that humans lack non-derived content, and that there is no such thing as original intentionality. Furthermore, in opposition to the view that cognitive content is non-derived, he noted that organisms have the minds they do in virtue of evolution by natural selection. This means that there is a kind of derivation of human minds, hence, Dennett believes, a kind of derivation of the content in their minds. Relating this to the debate over extended cognition, we can see that if all content is derived, then the content in Inga's brain is derived, just as is the content in Otto's notebook. So, the derived/non-derived distinction is irrelevant.

Clark (2005) has inspired a different line of attack on the need for cognitive states to bear non-derived content. He invites us to imagine the discovery of a Martian whose brain

manipulates bitmapped images of texts, so that the Martian would be manipulating texts in its brain. Clark claims that we would not hesitate to say that the Martian is doing cognitive processing with items bearing only non-derived content, so that we can see that cognitive processing does not necessarily traffic in only non-derived representations.[5] Clark (2010) brings the Martian example down to earth by proposing that certain human "memory masters" have the ability to recall image-like representations of texts. Were they to use their image-like representations on texts, then they would be performing cognitive operations on derived representations. Menary (2010) pursues this further by offering cases of "thinking in English." Words of natural language are paradigms of items bearing derived content, so insofar as "thinking in English" is literally thinking in English, we have cases of cognitive processing that involves trafficking in derived content.

Fisher (2009) argues that it is a mistake to propose that cognitive content must be non-derived, since there are numerous plausible cases of derived cognitive content. He suggests that the contents of thoughts, memories, and imaginational states derive from the contents of perceptual states. Given that the condition cannot hold of Inga's informational resources, it is too much to demand that it hold of Otto's.[6]

Substantial resistance from cognitive psychology?

A second potential source of substantial resistance to equating Inga and Otto comes from cognitive psychology. Rupert (2004) draws attention to two features of human memory: the phenomenon of negative transfer and the generation effect.[7] In the first phase of a typical negative transfer experiment, participants are presented pairs of items, such as a male name (*A*) paired with the name of their significant other (*B*), and trained to a criterion to respond with a *B* item when probed by an *A* item. In the second phase of the experiment, participants are told that the relations of all the *A* individuals have changed. Now, John is not married to Mary, but is married to Sue, and so forth. Thus, participants must learn a new set of *A-C* associations that replace the prior *A-B* associations. In this second phase, the new *A-C* associations are more difficult to learn that the initial *A-B* associations. This is the negative transfer effect. *Ex hypothesi* Inga will perform as a typical participant in one of these experiments. By contrast, Otto using his notebook will probably not. In the first phase, Otto will write down all the *A-B* pairings on a single page. This will, presumably, enable Otto to reach criterion on the first pass, performance that is quite different from Inga's. In the second phase, Otto will presumably write all the *A-C* pairings on a new page. This will enable Otto to reach criterion again on the first pass, thereby differing dramatically from Inga's performance. Finally, we see that, with Otto, there is no negative transfer. Learning the *A-B* list does not inhibit learning of the *A-C* list.[8]

In a typical generation effect experiment, there are two groups of participants. In the first group, participants are given a list of *A-B* pairs to memorize; in the second group, participants generate a list of *A-B* pairs by reading a sentence, such as "The cow chased the ball." In both groups, participants develop a list of *A-B* pairs, such as "cow-ball," but participants who generate the pairs by reading sentences have better memory recall. This again applies straightforwardly to the Inga-Otto case. If we suppose that Inga is a normal human being, then *ex hypothesi* she will perform better when she is in the group that generates its *A-B* list by sentence reading than when she is in the group that is given the *A-B* list. By contrast, whichever group Otto is in, he will write down the items in the *A-B* list, then recall as necessary later. He will perform just as well in the generation condition as in the non-generation condition.

The generation effect and negative transfer effect seem to show that Clark and Chalmers are mistaken in their contention that Inga and Otto are in all important and relevant respects the

same. Notice, as well, that this line of response can go beyond merely showing that Clark and Chalmers overstated their case for the similarity between Inga and Otto. We can draw attention to many other features of memory to argue, as does Rupert (2004), that external "memory" processing differs so dramatically from internal memory processing that these should not be treated as instances of a single natural kind. In advancing this argument, we do not have to say that any one feature, such as the negative transfer effect or the generation effect, is a necessary feature of memory. As noted above, we may, instead, propose that the particular features of memory that cognitive psychologists have so far discovered are elements in a much larger, perhaps poorly understood, homeostatic property cluster that constitutes human long-term memory as a natural kind. Thus, the more Inga and Otto differ in the properties that scientists have come to associate with long-term memory, the stronger is the case for saying that Inga and Otto do not have the same kind of "memory."

As emphasized above, Clark and Chalmers apparently believed that Inga and Otto are in *all* important and relevant respects the same.[9] Nevertheless, advocates of extended cognition need not take such a strong position. They can instead propose that Inga and Otto are in *some* important and relevant respects the same. If that were the case, there could be some sort of cognition that extends. This, in fact, seems to have become the favored defense of extended cognition.

So, what kind of important and relevant cognition is extended? There is no single answer. One proposal from Clark (2009) is that something like *folk psychology* is extended. Folk psychology consists of a largely implicit theory people have about the way humans perform in various situations. Perhaps, therefore, we could say that the hypothesis of extended *folk* cognition is true, important, and relevant, while conceding that the hypothesis of extended *scientific* cognition is false. Rowlands (2009) thinks that some form of *liberal functionalism* fits the bill.

> It is generally accepted that the arguments for EM [extended mind] presuppose functionalism. More than that, they presuppose a peculiarly *liberal* form of functionalism. Indeed, there is a way of understanding functionalism according to which EM emerges as a straightforward, almost *trivial*, consequence. In its more liberal forms, functionalism is based on a *principled* indifference to the details of the physical structures that realize mental processes. What is crucial to a mental state or process is its functional role, not its physical realization. For the liberal functionalist: if it walks like a duck, and talks like a duck, then it is a duck. *How* it manages to walk and talk like a duck is not directly relevant. To this, EM simply adds: neither does it matter where it walks and talks like a duck.
>
> *(Rowlands, 2009, pp. 56–57)*

Following Rowlands, then, we may say that the hypothesis of *liberal functionalist* cognition is true, important, and relevant, while conceding that the hypothesis of *scientific* cognition is false. Of course, whatever version of this approach we adopt, be it Clark's, or Rowlands', or some other, we do want some explanation of the importance of these forms of cognition and their extension and what relevance they have to scientific and philosophical projects.

Coupling arguments

Far more prevalent than cognitive equivalence arguments are a group of so-called "coupling arguments." So, to return to Clark and Chalmers' discussion, they argue that:

> It is not just the presence of advanced external computing resources which raises the issue, but rather the general tendency of human reasoners to lean heavily on environmental supports. Thus consider the use of pen and paper to perform long multiplication … , the use of physical re-arrangements of letter tiles to prompt word recall in Scrabble … , the use of instruments such as the nautical slide rule … , and the general paraphernalia of language, books, diagrams, and culture.
>
> *(Clark and Chalmers, 1998, p. 8)*

Then, somewhat later, they write

> In these cases, the human organism is linked with an external entity in a two-way interaction, creating a *coupled system* that can be seen as a cognitive system in its own right. All the components in the system play an active causal role, and they jointly govern behavior in the same sort of way that cognition usually does. If we remove the external component the system's behavioral competence will drop, just as it would if we removed part of its brain. Our thesis is that this sort of coupled process counts equally well as a cognitive process, whether or not it is wholly in the head.
>
> *(Ibid., pp. 8–9)*

These passages invite the following interpretation: when a process X has a causal influence on a cognitive process Y, then we may infer that the entire X-Y process is a cognitive process. This interpretation is reinforced if we bear in mind the introduction to Clark and Chalmers (1998), where they write, "We advocate a very different sort of externalism: an *active externalism*, based on the active role of the environment in driving cognitive processes" (ibid., p. 7).

While there are texts that suggest this line of argumentation, it is now common ground among advocates and critics of extended cognition that this simple argument is fallacious.[10] Just because a process X causally influences a cognitive process Y does not seem to be a sound basis upon which to infer that the whole X-Y process is cognitive. Consider this kind of reasoning in another context. Let X be the process of evaporative cooling found in a typical home air conditioning system. In a normal working system, this process is causally coupled to a process Y of deformation in the shape of a bimetallic strip in a thermostat in one room of the house. This causal linkage between X and Y does not make the entire X-Y process a process of shape deformation. For that matter, it does not make the whole X-Y process a matter of evaporative cooling.

While many advocates of extended cognition admit that there is this "coupling-constitution fallacy," they also sometimes suggest that this is not a mistake anyone has committed and leave the matter at that. Alternatively, they take the argument to be that cognition extends when there is the *right kind* of causal connection between brainy cognitive processes and bodily and environmental processes. The simplistic coupling-constitution fallacy sketched above does not speak to this more subtle claim.

These replies, however, underestimate the challenge of moving from causal claims to constitutive claims. Rather than viewing the simple coupling-constitution fallacy as the end of a story, we should take it to be a starting point. It should lead us to ask how to refine the mistaken, simplistic account into a correct, sophisticated account. There are, however, two challenges lurking here. First of all, if we propose that additional conditions, C_1, … , C_n, are needed in order for coupling relations to suffice to establish constitutive relations, then we must bear the burden of determining and theoretically validating those additional conditions C_1, … , C_n. Second, once we establish a valid set of conditions, if we subsequently wish to show that we

have an actual instance of extended cognition, as when using pen and paper to perform a long multiplication, then we will have to show that the putative actual instance satisfies C_1, … , C_n. Moreover, the challenges interact. The more restrictive the set of theoretically valid conditions we develop, the greater the argumentative burden that will be needed in order to establish that they apply in any particular putative instance. Moreover, the more restrictive the set of theoretically valid conditions we develop, the fewer real-world instances we are likely to find to satisfy them.

To get a sense of this, let us consider a relatively trivial strengthening of the causal conditions that might be thought to suffice for a constitutive relation. Perhaps we should say that when a process *X* has a *reliable* causal influence on a cognitive process *Y*, then we may legitimately infer that the entire *X*-*Y* process is cognitive. Clark and Chalmers seem to have such an idea in mind when they write,

> for coupled systems to be relevant to the core of cognition, *reliable* coupling is required. It happens that most reliable coupling takes place within the brain, but there can easily be reliable coupling with the environment as well. If the resources of my calculator or my Filofax are always there when I need them, then they are coupled with me as reliably as we need. … If the relevant capacities are generally there when they are required, this is coupling enough.
>
> *(Clark and Chalmers, 1998, p. 8)*

Now return to the two challenges. In the first place, is a reliable causal connection really enough to warrant the conclusion that there is an extended cognitive process? Revisit the example of an air conditioning system. Let the process *X* of evaporative cooling be *reliably* causally coupled to the process *Y* of deformation in the shape of a bimetallic strip. This does not seem to make a difference. This reliable causal linkage between *X* and *Y* does not make the entire *X*-*Y* process into a process of shape deformation. In this instance, reliable causal connection does not seem to determine the character of a process. So, the addition of the reliability condition does not seem to suffice to circumvent the coupling-constitution fallacy.[11] In the second place, reflect on the implications of the reliability condition. With this condition, not all uses of pen and paper for long multiplication will be instances of extended cognition. Cases in which an agent always has pen and paper at hand to do long multiplication will count as instances of extended cognition, where cases in which those implements are only occasionally available will not. This means that there are likely to be fewer instances of extended cognition than Clark and Chalmers may have led us to expect.

For a second illustration of the interaction between the two challenges, we could consider a set of more stringent conditions proposed by Clark and Chalmers:

> First, the notebook is a constant in Otto's life—in cases where the information in the notebook would be relevant, he will rarely take action without consulting it. Second, the information in the notebook is directly available without difficulty. Third, upon retrieving information from the notebook he automatically endorses it. Fourth, the information in the notebook has been consciously endorsed at some point in the past, and indeed is there as a consequence of this endorsement.
>
> *(Clark and Chalmers, 1998, p. 11)*

One may think that these conditions adequately meet the first challenge by yielding correct results for cases such as Otto's use of his notebook, but they also appear to rule out instances of

bona fide memory. They appear to disallow cases of "alienation" from one's own cognitive resources. So, suppose that, for whatever reason, Dotto comes to the conclusion that he should not trust his memory to correctly recall people's names. Although he always does correctly recall the name of a person he sees, he does not trust his memory. Rather than take a chance on misidentifying someone, he either avoids the person or asks someone nearby to confirm his recollection. In this scenario, Dotto violates the first and third of Clark and Chalmers' conditions, suggesting that Dotto's memory is not really involved in cognitive processing.[12]

These sorts of examples, of course, do not prove that it is impossible to meet the two challenges. They do not show that it is impossible to find actual cases in which cognition extends when there is the *right kind* of causal connection between internal cognitive processes and bodily and environmental processes. Instead, they evidently establish the coupling-constitution fallacy as a provocation to further investigation.[13]

Conclusions

This chapter has focused on two prominent types of arguments for extended cognition, cognitive equivalence arguments and coupling arguments. These are perhaps best thought of as families of arguments. The members of the cognitive equivalence family differ among themselves regarding the standard of cognitive equivalence – for example, equivalence in all important and relevant respects, or equivalence in one or another respect. The members of the coupling family differ among themselves regarding what conditions, over and above a simple causal connection, warrant the conclusion that some process external to the brain constitutes a cognitive process. Insofar as we are concerned with the conditions under which there might be extended cognition, these appear to be the core approaches.[14]

Notes

1 The only other well-known cognitive equivalence argument is Clark and Chalmers' thought experiment with the video game Tetris. This argument was inspired by Kirsh and Maglio (1995).
2 In a similar bold spirit, Mark Rowlands claims "there is no sound theoretical reason for setting up a dichotomy between internal memory processes and external aids to those processes" (Rowlands, 1999, p. 121).
3 See, for example, Adams and Aizawa, 2001, 2008; and Fodor, 2009.
4 See, for example, Boyd, 1999a, 1999b.
5 For further discussion, see Clark, 2010.
6 For further discussion, see Adams and Aizawa, 2008; and Aizawa, 2010.
7 See Adams and Aizawa, 2008 for other examples of this sort.
8 Of course, we might try to modify the original Otto example in such as way as to display the negative transfer effect. Perhaps Otto does not write the *A-C* pairings on a new sheet of paper, but imperfectly erases the *B* entries and writes the *C* entries where the old *B* entries were. This way of entering the pairings may make it harder for Otto to write them correctly on the first try. And maybe this will make it harder for Otto to read the *A-C* pairings, which might explain why it takes him longer to get the *A-C* pairings right.
9 Also recall Rowlands' claim that "there is no sound theoretical reason for setting up a dichotomy between internal memory processes and external aids to those processes" (1999, p. 121).
10 See, for example, Clark, 2009; Fisher, 2009; and Wilson, 2010.
11 Notice that the reciprocal causal interaction between *X* and *Y* does not appear to make a difference. And, would it make a difference if the interaction between *X* and *Y* were non-linear?
12 For other examples in this vein, see Adams and Aizawa, 2008, ch. 7.
13 Other sets of condition on coupling will be found in van Gelder (1995) and Haugeland (1998). More recently, Chemero (2009) and Froese, Gershenson, and Rosenblueth (2013) have proposed that cognition extends when there is a non-linear coupling between brain processes and environmental

processes, where Roberts (2012) has proposed that true cognitive extension occurs only when the subject takes responsibility for the contribution made by a non-neural resource.

14 For an alternative to causal coupling arguments, see Gillett, 2013, and the commentary in Aizawa, 2013.

References

Adams, F., and Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology, 14*(1), 43–64.

——(2008). *The bounds of cognition*. Malden, MA: Blackwell Publishers.

Aizawa, K. (2010). The boundaries still stand: A reply to Fisher. *Journal of Mind and Behavior, 31*(1–2), 37–47.

——(2013). Introduction to "The material bases of cognition." *Minds and Machines, 23*(3), 277–86.

Boyd, R. (1999a). Homeostasis, species, and higher taxa. In R. Wilson (Ed.), *Species: New interdisciplinary essays* (pp. 141–85). Cambridge, MA: MIT Press.

——(1999b). Kinds, complexity and multiple realization: Comments on Millikan's "Historical kinds and the special sciences." *Philosophical Studies, 95*(1), 67–98.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.

Clark, A. (2005). Intrinsic content, active memory and the extended mind. *Analysis, 65*(1), 1–11.

——(2009). *Supersizing the mind: Embodiment, action and cognitive extension*. New York: Oxford University Press.

——(2010). Memento's revenge: The extended mind, extended. In R. Menary (Ed.), *The extended mind* (pp. 43–66). Cambridge, MA: MIT Press.

Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19.

Dennett, D. C. (1986). The myth of original intentionality. In K. A. Mohyeldin Said, W. H. Newton-Smith, R. Viale, and K. V. Wilkes. (Eds.), *Modeling the mind* (pp. 43–62). Oxford: Oxford University Press.

Fisher, J. (2009). Critical notice of the book *The bounds of cognition*, by F. Adams and K. Aizawa. *Journal of Mind and Brain, 29*(4), 345–57.

Fodor, J. (2009). Where is my mind? *London Review of Books, 31*(2), 13–15.

Froese, T., Gershenson, C., and Rosenblueth, D. A. (2013). The dynamically extended mind: A minimal modeling case study. Paper presented at the IEEE Congress on Evolutionary Computation, CEC 2013, Evolutionary Robotics track, Cancun, Mexico, June.

Gillett, C. (2013). Constitution, and multiple constitution, in the sciences: Using the neuron to construct a starting framework. *Minds and Machines, 23*(3), 309–37.

Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Ed.), *Having thought: Essays in the metaphysics of mind* (pp. 207–37). Cambridge: Harvard University Press.

Kirsh, D., and Maglio, P. (1995). Some epistemic benefits of action: Tetris, a case study. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Menary, R. (2010). The holy grail of cognitivism: A response to Adams and Aizawa. *Phenomenology and the Cognitive Sciences, 9*(4), 605–18.

Roberts, T. (2012). You do the maths: Rules, extension, and cognitive responsibility. *Philosophical Explorations, 15*(2), 133–45.

Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*, New York: Cambridge University Press.

——(2009). Enactivism and the extended mind. *Topoi, 28*(1), 53–62.

Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy, 101*(8), 389–428.

van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy, 92*, 345–381.

Wilson, R. (2010). Extended vision. In N. Gangopadhyay, M. Madary, and F. Spicer (Eds.), *Perception, action and consciousness: Sensorimotor dynamics and two visual systems* (pp. 277–90). Oxford: Oxford University Press.