My BlackBerry and Me: Forever One or Just Friends?

Andrew Brook

Institute of Cognitive Science

Carleton University

Ottawa, Canada

Abstract

The idea that something crucial to cognition lies outside the brain and even the skin is central to a number of significant recent developments in cognitive science and philosophy of mind. Loosely grouped under the label 'situated cognition', all variants of the view hold that the context of cognition is crucial to it. However, there is a big division over in what way. Indeed, there are a number of divisions over what way, but the one on which I will focus is this: Are contexts just causally linked to cognition or are they actually part of cognition, some contexts anyway? However crucial researchers of the first persuasion take contexts to be, they do not mess around with the traditional boundary between mind and world. Researchers of the second persuasion do challenge that boundary. They hold that in some respects cognition actually extends beyond the brain and skin, that something outside the body is part of cognition. This idea can take at least three forms:

(i) A claim that there is a normative and therefore a social dimension to cognition (the normativity of mind hypothesis),

(ii) A claim that semantic and representational content consists of a relationship between representational vehicles and something else (externalism).

(iii) The hypothesis of the extended mind (Clark and Chalmers 1998).

After a survey of situated cognition and some comments on the normativity view, we will zero in on externalism and the extended mind hypothesis. The internet and wireless appliances such as the BlackBerry are an interesting test case, especially for the latter. On some views, my BlackBerry is part of the cognitive system that is me. On a credible view of the boundaries of the mind, there is little or nothing to support such a view.

1. "Cognition ain't (all) in the head" (Clark and Chalmers)

The idea that something about cognition lies outside the brain or skin is central to a number of significant recent developments in cognitive science and philosophy of mind, many loosely grouped together under the name, 'situated cognition'. Mobile devices such as the Blackberry and even the internet in general have been thought by some to support versions of the view, even quite radical versions. And some think that the idea of situated cognition will help us understand what mobile media and the net are doing to us.

There is a cynical saying about philosophy. Either a philosophical idea is interesting, in which case it is almost certain to be false, or it is true, in which case it is almost certain to be uninteresting. (The same has been said about cognitive science.) Situated cognition is certainly interesting. Does any version of it have a chance of being true?

The term 'situated cognition' could be applied to every view that we will discuss but it is often reserved for what are in fact the tamest expressions of the view, positions such as those advocated by Ed Hutchins, David Kirsh, Rodney Brooks, the Andy Clark of *Being There,* and the like. These are certainly the best known variants of the view. The core idea behind all of this group of positions is that to understand cognition, we have to study it in the contexts in which it occurs, physical, social, and historical. It cannot be studied effectively in isolation.

Situated cognition of this sort comes in many flavours. In rough order of strength, there is:

Distributed Cognition – Cognition is "stretched over, not divided among, mind, body, activity, and culturally organized settings" (O'Connor and Glenberg 2003)

Ecological Psychology – We need to pay attention to the environments of cognition and how cognition situates itself in them.

Embedded Cognition – Cognition is embedded in physical system(s).

Embodied Cognition – Cognition takes place in a body and this shapes it significantly.[1]

Some key claims characterize many theories in this group:

There is much less representing of the world going on in us than classical cognitive science would suggest (Brooks 1991, Clark 1997).[2]

In cognitive activity, the world is often its own model (Brooks 1991).

Cognition relies massively on 'cognitive artefacts' (Hutchins 1995, Clark 2001).

Perception is largely a grand illusion (O'Regan, Noë) and action, not representation, is the heart of cognition (the enactive approach of Varela, Thompson, and Rosch, Noë, and so on.)

Cognition is largely (for some theorists entirely) evolution-picked special-purpose modules (Barkow, Cosmides and Tooby 1992, Pinker 1997).

Temporal dynamics are important; in fact, dynamic system theory (DST) is apt to be a better guide to cognition than logic and classical AI (Port and van Gelder 1995).

Notice something interesting: However exotic these claims may be, nothing in any of them requires that we mess around with the traditional boundaries of the mind. So much for our quick overview of situated cognition.

The next group of views to be discussed holds that there is a normative and therefore a social and/or evolutionary dimension to mindedness. To be a mind is to perform certain assigned functions properly and/or to satisfy other norms. This notion has been espoused mainly by philosophers. Davidson, Dennett, Millikan, Dretske, Brandom, and Haugeland are leading figures.

The claim that mindedness has a normative dimension is stronger in some ways than the claim that cognition is situated. The normativity hypothesis has to insist upon something about which situated cognition can be neutral: that something outside brain and skin is required, is necessary, for anything to be a mind. Nonetheless, they are alike in one respect. They can both accept that the

---

[1.] In turn, the idea of embodiment comes in many different flavours (Chrisley and Ziemke 2003): Physical realization; physical embodiment; realization in an integrated system; organismoid embodiment (some body is needed); organismal embodiment (an organic body much likes ours is needed); and maybe some others.

[2.] Some people in the situated cognition movement have even claimed that no cognition does not contain representations. Clark (1997) shows that this is at best an overstatement. Representations cannot be eliminated entirely.

boundaries of the mind do not or need not extend beyond the brain or at any rate the skin.[3] An external element is necessary for mindedness, but it need not be part of what a mind is.

There are at least two views which hold that something external to brain and skin is part of what a mind is. We will devote the rest of the paper to them. One is semantic or representational externalism, the idea, to use Putnam's (1975) famous phrase, that 'meaning ain't in the head'. Call this the hypothesis of extended content (EC). The other is Clark and Chalmers' (1998) hypothesis of the extended mind (EM).

Semantic or representational externalism maintains that part of what gives meaning to syntactically structured strings, part of what gives representational content to representational vehicles, consists of a relationship between structures and vehicles and something else. There is a wide variety of candidates for what this something else consists in:

Causal relationships that reliably result in tokenings of the appropriate vehicle

Nomological relationship that do the same (both views advanced by Fodor 1994)

The function of representing something (Dretske 1995, Millikan 1984)

Experts or some other social factor (Burge and a host of subsequent philosophers)

What these theorists all agree upon is that for strings of words to have meaning, for representational vehicles to have content, the string or vehicle must be connected to something outside the string or vehicle and also outside the brain or skin.[4]

Now Clark and Chalmers' (1998) EM. Where externalists maintain that part of *content* is beyond brain and skin, EM maintains that *the mind itself* extends beyond brain and skin.

> If, 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 part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain't (all) in the head!

There is a way to make this claim trivially true: stretch the term 'cognitive system' to cover mind *plus* the things that it makes use of in its cognitive activities. This would be something like taking cognition *plus* context and relabelling the whole thing 'cognition'. You could do this if you wanted but it would gain you nothing. Clark and Chalmers make clear, however, that they have something more radical than this in mind. Their claim is that *the mind itself* extends beyond brain and skin. Hence their title. Up to now, we have switched back and forth between 'cognitive system' and 'mind' rather unreflectively. From now on, we will try to use these terms with precision, to reflect the distinction just made.

2. Does EC support EM?

Does externalism (EC) support EM (extended mind hypothesis)? The *ur*-thought experiment for nearly all forms of externalism is Putnam's (1975) twin earth thought experiment (TETE). Here is how the story goes. Suppose that Adam and his twin on twin earth, Twadam, have beliefs about a certain clear liquid in front of them and both call it 'water'. One liquid is H^2O and one is XYZ. If beliefs are individuated by what they are about, by their content, it would appear that Adam and Twadam have different beliefs. However, the element of content that makes the beliefs different is

---

[3] This is not to say that all these theorists do accept that.

[4] When assigned function is the external element (Dretske, Millikan), it may be external to the representational vehicle but it is not clear that it is external to the brain or skin.

external to them (indeed, it is not even known to them). So part of what makes their beliefs the beliefs that they are is external to the two believers.

It is not widely known but TETEs can be used as arguments for the other forms of externalism, too. What Adam and Twadam are causally related to varies. If so, the content-supporting nomological laws would vary. Moreover, the function assigned to the string 'water' could vary. Finally, experts' knowledge could vary. All these while the internal states of the two people remain the same. TETE is an all-purpose externalism generator.

Does TETE work? As Segal (2000) has pointed out, it is not clear that the external element does in fact affect the content of the two beliefs. Adam and Twadam both believe what could be expressed as, "This [the substance in front of them] is water." However, if they don't also believe, "This is water-rather-than-twin-water" (in their respective idiolects), then their concept of water in each case may well be broad enough to range over both substances. If so, their beliefs would have the same content. (Brook and Stainton 1997 reach the same conclusion in a different way.[5])

Even if TETE were to work, what could it show? Not as much as some have thought. The most it could show is that EC, something in the *content* of beliefs, representations, and so on, is external to the brain or body. It would not show EM, that *the mind itself* is extended. In fact, externalism says nothing about the boundaries of the mind itself.

3. Clark and Chalmers' argument for EM

The argument that Clark and Chalmers offer for EM is very simple. EC does not go far enough because in it the element external to brain and skin is cognitively inert. There are also, they urge, processes external to brain and skin that are cognitively active. Then, as we saw, they claim that, were these processes carried out in the head, they would be part of the mind. So if I use a pad of paper and pencil, or a calculator, to do arithmetic, that pad of paper, that calculator, are part of the mind that I am. In short, EM draws on semantic or representational externalism – and also on the situated approach sketched earlier – but is more radical than either of them.

4. My BlackBerry, the Web, and EM

BlackBerries and the Web are a nice test cases for EM.[6] They both put pressure on EM, in at least two ways.

1. Like situated cognition generally, EM thinks of the brain (the internal part of the mind, on their view) being coupled to the world in much the way a control system with feedback is coupled to the system it controls. Clark (1997) calls it *continuous reciprocal causation.* With wireless devices and the internet, coupling to the brain is doubtless causal but it is a causal process very different from reciprocal causation.

There is a role for standard coupling – to create a communication channel between the brain and some electronic device. Once this channel is in place, however, standard coupling is done. What happens next is very different. What happens next is *semantic* coupling between brain and device. In

[5] Externalism has been thought to create problems for the relationship of conscious content, which does seem to be in the head, to representational content and also for first person authority. Myself, I don't see a problem, but that is a story for another day.

[6] I choose the Blackberry for patriotic reasons: RIM, their maker, is a Canadian company.

fact, so long as a communication channel for semantic coupling has been set up, the nature of the channel and the connected device (Blackberry or PalmPilot or Nokia, Dell or HP or Compaq) is completely irrelevant.

Semantic coupling is very unlike control system coupling. It is a matter of co-reference, implication ('if this, then that', 'this is inconsistent with that'), a shared syntax, shared conventions about meaning, and so on. In addition, it has normative dimensions – a presumption of truth-telling, of an intention to communicate, and so on.[7]

What semantic coupling is actually like puts one kind of pressure on EM. If semantic coupling is so different from continuous reciprocal or any other form of standard causal coupling, what could be meant by saying that this relationship constitutes an *extension* of the mind beyond the brain? In semantic coupling, the relation of what is in the head or, for the matter, what is not in the head to the whole is not a part/whole relationship of any usual sort, so EM is not using the word 'extended' in any usual way.

2. Next, notice the nature of the extra-brain element into which mind as said by EM to extend. The extended element of *my* mind will often be *someone else's* mind. So if EM is sound, everyone I know, no, everyone connected to me via the web in any way could be part of my mind! Now, I don't object to this – I've always thought that I had pretty big mind. (Others substitute 'head'.) Trouble is, everyone reading this gets to say exactly the same thing! And I don't know how to think of me being a part of .... a part of .... me.

More generally, we end up with bizarre picture of humankind. Instead of the six-billion-odd separate minds that we thought we were, every mind would be populated by millions of other minds fading out in thinner and thinner semantic relations. Imagine the resulting Venn diagramme! A very messy picture.[8]

Indeed, the EM picture is too messy to be tolerated if we can find a credible alternative. I think we can. First, recall the distinction between minds and cognitive systems. Consider again the Clark/Chalmers criterion, whether a process would be considered part of the mind if it were done in the head. We could make it a matter of definition that cognitive systems extend beyond the head; the Clark/Chalmers criterion may even work for deciding the extent of a cognitive system. Minds are a different case, however. There are some persuasive reasons for saying that minds do not extend nearly as far.

How far *my mind* extends, it is plausible to hold, starts from how far *I* extend and then adds elements as necessary. Note that the relevant notion of me is a *forensic concept*, as Locke (1690) put it. I, the mind that I am, is the locus of agency, the thing *responsible for my actions* (Preston, this volume). So how far do I extend?

So how far do I extend? One natural answer is: as far as my skin. As Clark and others insist, there is no magic to the skin but it does mark some important *functional* boundaries. If I close one

---

[7] How these work out in the context of the internet might be an interesting topic in its own right. Note that there is a strange bifurcation in cognitive science between philosophers and experimental psychologists on the topic of semantics. Philosophers have a huge problem taming it in a naturalistic account. Some even claim that the brain can't do semantics and then frantically search around for a surrogate. (Dennett 1987, Ch. 4, p. 61 and Fodor 1994, Ch. 1, pp. 8-14 are two examples.) Meanwhile, psychologists blithely assume that their subjects can process semantic information just fine and get on with their experiments.

[8] The problem extends to situated cognition, too. Andy Clark says that he is willing to bite this bullet (private communication).

pair of eyelids, I stop seeing. If I prick one system of skin, I feel pain. No other eyelids or skin do these things for me. So one natural criterion for how far I extend is C1:

> **C1.** I extend to the periphery of the body on which I am perceptually and sensationally dependent.[9]

Another natural answer is that I extend as far as the behaviour I can bring about by willing a basic action, i.e., an action that I can bring about without intentionally bringing about any other action. So,

> **C2.** I extend as far as the movements I can control by willing basic actions,

Neither C1 nor C2 ends the mind *neatly* at the skin. Consider prosthetic devices and tools. We can be dependent on them for perceptions and, arguably at least, they can be controlled by basic actions. That we end at the skin most of the time is a functional matter. However, the isomorphism is not too bad most of the time.

Another intuitively-appealing idea of how far I extend is that I extend as far as my mental states extent. Such states are not just related to me. They are part of me. A tooth-ache I feel is part of me. A tooth-ache you feel is not. An act of imagining that I do is part of me. An act of imagining that you do is not. And so on. And how do we delineate which mental states are mine? Here is one way:

> **MS_{mine}** A mental states is mine if I can access it 'from the inside.' This metaphor can then be cashed in two ways:
>
> > **MS_{mine1}** I access the state by having, feeling or doing it.
> >
> > **MS_{mine2}** I access the state from the point of view of having, feeling or doing it.

On MS_{mine}, when I access an experience by having it, an emotion by feeling it, an action by doing it, the experience, emotion, or action is part of me.[10] Note that **MS_{mine2}** is broader than **MS_{mine1}**. When I *remember* having an experience, I am not accessing it by having it but I am accessing from the point of view of having it. We can now state C3:

> **C3:** I extend over the range of states that I can access by or from the point of view of having, feeling, or doing them.

Because my access to my bodily states via proprioception is from the point of view of having, feeling, or doing them and proprioception extends to the skin, C1 and C3 draw the boundary around me in much the same place, even though they are quite independent of one another.

Finally, picking up Locke's idea that the concept of a person (in our terms, a mind) is a forensic concept, we might fix my extent by reference to states of affairs that can be part of my processes of focussed, intentional decision-making. Here, however, we have to be careful. Since I can take any state of affairs into account in making my decisions, we have to distinguish between states that I can *take into account* and states that are *part of the decision-making process.* This distinction suggests C4:

> **C4:** I extend over the states that are part of my decision-making processes.

Here is one way to make the underlying distinction. States that are part of my decision-making process get to be so simply by my consciously having them. I do not need to represent them in any other way. For example, for a desire to enter into my making a decision, I must be aware of it. But to

---

[9.] P. F. Strawson explored a similar idea in the 'Persons' chapter of (1959), p. 90-1.

[10.] This criterion depends on there being self-presenting representations and thus entails that one does not need a higher-order thought to be conscious of a toothache or a desire (*pace* Rosenthal 1993) and one does not become conscious of one's states by inference from what the states make one conscious of (*pace* Dretske 1995).

be aware of it, it is enough that I *feel the desire.* I do not need to represent it in any other way. However, for the state of the weather to enter into my making a decision, it is not enough that it exist. I must also represent it.

A great deal more needs to be said here, about the unity of consciousness and the nature of the agent of unified consciousness in particular. However, even with what little we have said, something interesting has happened. The criterion for what is part of decision-making in connection with C4 is one of the two that we used to delineate the states relevant to C3. If so, on this criterion they have to draw the boundary in the same place. That place is again, roughly, the skin.

We now have a reasonably plausible account of how far *I* extend (and therefore, on the broad concept of a cognitive system introduced earlier, of what part of a cognitive system *I* would make up). On anybody's account, however, the mind that I am extends farther than I do. One extension is obvious. The states that I *could* access by or from the point of view of having, feeling, or doing them are just as much part of me as the ones that I *actually am* thus accessing at any given moment. It would be implausible (though not impossible – see below) to deny that my memories are part of me, for example, yet at any one time I am accessing only a tiny fraction of them. Limiting me to what I am actually accessing would yield a small me indeed – though I would then be a perfect Cartesian mind.

'Implausible though not impossible': Clark and Chalmers mount an extended comparison between memory and external prosthetics such as notebooks. They conclude that any differences between their Inga, who has intact memory, and Otto, whose memory is severely impaired and who has to rely on notebooks, are 'shallow'. They then argue that since memory is part of the mind, the prosthetics are, too. Well, one person's modus ponens can another person's modus pollens: one could just as rationally *exclude* both of them. But there is a better way to go.

While Otto's notebooks play an indispensable role in the cognitive system of which he is a part, and therefore *for* him, they are not *part of* him. There are at least two important difference between them and memories. First, Otto can access his (sorely diminished store of) memories just by calling them up, and many of them will be memories of having, feeling, and doing various things. To access his notebooks, however, Otto has to perceive them – i.e., represent them in states or events that are different from the states or events that make up the books. Second, when he does look at his notebooks he does not access the original states and events from the point of view of having them. A description of having done something in a notebook, no matter how detailed, would not allow one to re-experience the event in the way that remembering the event does – or at least can.

Are there states outside my skin that are or could be part of me? We have already said that prosthetics and tools may create such states. If one person could share two bodies, that would be another way. To return to where we began, however, my BlackBerry is not part of me. And neither is the internet. Whew![11]

References

Barkow, J. H., Leda Cosmides & John Tooby. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture.* New York: Oxford University Press.

Brook, A. and R. Stainton. 1997. Fodor's New Theory of Content and Computation, *Mind and Language* 12, No. 3 & 4, pp. 459-74

[11.] Thanks to Zoltán Jakab and Andy Clark for helpful conversation, to John Preston for allowing me see a draft of his fine paper, and to Kristof Nyiri for inviting me to write this paper.

Brooks, R. 1991. Intelligence without representation. *Artificial Intelligence* 49, pp. 139-159

Clark, A. 1997. *Being there.* Cambridge, MA: MIT Press

Clark, A. 2001. *Mindware.* Oxford: Oxford University Press

Clark, A. and D. Chalmers. 1998. The extended mind. *Analysis* 58, pp. 10-23

Chrisley, R. and T. Ziemke. 2003. Embodiment. Lynn Nadel, ed. *Encyclopaedia of Cognitive Science*. London: Nature Publishing Group, pp. 1102-1108

Dennett, D. 1987. *The Intentional Stance.* Cambridge, MA: MIT Press

Dretske, F. 1995. *Naturalizing the Mind.* Cambridge, MA: MIT Press

Fodor, J. 1994. *The Elm and the Expert.* Cambridge, MA: MIT Press

Hutchins, E. 1995. *Cognition in the Wild.* Cambridge, MA: MIT Pres

Locke, J. 1690. *An Essay Concerning Human Understanding,* ed. and introd. P. H. Nidditch. Oxford : Clarendon Press, 1975.

Millikan, R. 1984. *Language, Thought, and Other Biological Categories.* Cambridge MA: MIT Press/A Bradford Book.

O'Connor, K. and A. Glenberg. 2003. Situated cognition. Lynn Nadel, ed. *Encyclopaedia of Cognitive Science*. London: Nature Publishing Group, pp. 19-25

Pinker, S. 1997. *How the Mind Works.* New York: Norton.

Port, R. F. and T. van Gelder, eds. 1995. *Mind as Motion: Explorations in the Dynamics of Cognition.* Cambridge, MA: MIT Press.

Preston, J. this volume. Is your mobile part of your mind?

Putnam, H. 1975. The meaning of 'meaning'. In his *Language, Mind and Reality* Philosophical Papers Vol. 2, pp. 215-71.

Rosenthal, D. 1991. *The Nature of Mind.* Oxford: Oxford University Press

Segal, Gabriel. 2000. *A Slim Book about Narrow Content.* Cambridge, MA: MIT Press

Strawson, P. F. 1959. *Individuals.* London: Methuen