

Katalin Farkas

Extended mental features¹

in: Matteo Colombo, Elizabeth Irvine, & Mog Stapleton (eds.) *Andy Clark & Critics* Oxford University Press 2019: pp. 44-55

Abstract: The focus of the original argument for the Extended Mind thesis was the case of beliefs. It may be asked what other types of mental features can be extended. Andy Clark has always held that consciousness cannot be extended. This paper revisits the question of extending consciousness.

Keywords: extended mind, extended cognition, functionalism, consciousness, perception

1. EM and the case of Otto

One of Andy Clark's main contributions to our understanding of the mind is the powerful and sophisticated defense of the so-called „extended mind” (EM) thesis. According to this thesis, the material basis of certain mental states can include, as constituents, parts of the world that are outside our organic body. Clark has always held that EM doesn't apply to consciousness. In this paper, I will briefly recall Clark's reasons for holding that the argument for EM cannot be used to support “ECM”, that is, Extended Conscious Mind. I will argue that the resistance to ECM is in tension with Clark's description of two cases which prima facie do support ECM: the example of an agent watching a computer screen while playing Tetris, and a problem-solving task in an experiment by Dana Ballard and co-workers. I propose that the tension should be resolved by reiterating the commitment against the extension of consciousness, but thoroughly revising the description of the Tetris and Ballard et al. cases. Finally, I will suggest a possible reformulation of the debate that hopefully puts the main issues into sharper focus.

The Extended Mind thesis was first explicitly defended in a joint paper by Clark and David Chalmers in 1998. As Clark reflected a few years later, “the original argument for EM took the form of some thought experiments, a principle meant to command rational assent, and a battery of responses to probable objections.” (Clark 2005 pp.1-2). One of the thought experiments was the subsequently much-discussed case of Otto, who suffers from memory loss and stores information in a notebook, and Inga, who, in contrast, has normally functioning memory. Suppose that both Inga's biological memory and Otto's notebook contain the information that the Museum of Modern Art is on 53rd street. Clark and Chalmers argued that – with the story suitably set up – the process of Otto

I would like to thank the editors of the volume for extremely helpful comments on an earlier draft of the paper.

retrieving information from his notebook is on a par with Inga's retrieving information from internal memory. If we think Inga has a belief with this content prior to recalling the information, Otto should be credited with the same belief. The fact that in Otto's case the information is stored outside his body is not relevant.

The rationally compelling principle mentioned by Clark has become known as the “Parity principle”, and goes as follows:

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. (Clark and Chalmers 1998 p.8)

The Parity Principle is supposed to help us to overcome our “bio-chauvinistic prejudices” (Clark 2005, p. 2) which locate the vehicles of mental processes within the body or the skull. By the Parity Principle, Otto's notebook is just as much part of a cognitive process as Inga's brain-states storing her memories. If this is not immediately compelling, Clark asks us to consider another case: a Martian whose biological memory is capable of storing bit-mapped images of texts. The Martians can store and recall these images on demand, just as we can store and recall information in our biological memory. “Surely, in that case”, says Clark, we “would have no hesitation in counting these goings-on as part and parcel of Martian cognition”. (Clark 2005 p. 5) But by the Parity Principle, we have to reach the same conclusion about Otto's notebook.

2. Against Extended Conscious Mind

The extended mental state that is attributed to Otto is a dispositional belief, a standing state that is not part of the stream of consciousness. This is important, because from the very beginning, Clark and Chalmers have expressed scepticism about the possibility of extending conscious features.

Clark argues at length that consciousness probably doesn't extend in a 2009 paper called “Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head”. He contrasts the original Extended Mind thesis which concerned “the vehicles of non-conscious mental states such as states of dispositional believing” (Clark 2009 p. 967) with a more radical thesis, “Extended Conscious Mind” (ECM), which extends the material basis of conscious experiences outside the brain.

Clark offers an illustration of the kind of extension that defenders of ECM promote. On the standard, non-extended picture, motion perception is realised in a certain part of the brain. Evidence

suggests that this is the Middle Temporal area, whose selective damage results in a disruption of motion perception (Marcar, Zihl, and Cowey 1997, referred to by Clark 2009 p. 968). In contrast, ECM suggests that the physical basis of at least some conscious experiences extends further into the non-neural body and the world. “It is this kind of claim that I want to examine” Clark continues.

I hope to show that nothing in the arguments for EM should incline us to accept ECM (to accept an extended view of the mechanisms of the conscious mind or of the vehicles of conscious experience). (Clark 2009 p. 968)

Clark believes that one of the strongest arguments for ECM comes from the dynamic entanglement of neural activity, bodily states and the world, in the processes of perception and action. However, just noting these extended dynamic loops is not sufficient to support ECM; what is needed in addition is that specifically the loops that support conscious experience run outside the brain. Clark thinks this is not *prima facie* plausible, “(f)or intuitively, visual experience may at each moment depend solely on the complex brain activity caused (instrumentally) by the ongoing engagements with body and world.” (Clark 2009, p. 977). To try to settle the issue, we could apply the following “twin test”: would a neural duplicate of yours, brought to existence by some quantum accident, share your experiential states? If the answer is yes, ECM lacks support; if the answer is no, ECM is vindicated. After some back and forth on the question of the temporal spread of experiences, Clark concludes that this particular dialectic ends in a stalemate: neither ECM, nor its opponent has a conclusive case.

However, a final set of considerations tips the balance against ECM. Clark proposes to develop a suggestion by Chalmers, that

consciousness requires direct access to information on an extremely high-bandwidth [...] our low-bandwidth conscious connection to the environment seems to have the wrong form as it stands. (Chalmers 2008, pp. xii–xiii, quoted by Clark 2009, p. 983)

This idea does not threaten the Otto case where there is no online task-solving and hence speed is not an issue. Clark hopes that this point, at the end, provides a principled reason to keep the original arguments for EM while resisting ECM. The upshot is that

we can allow that at least some conscious experiences may have the precise qualitative nature they do only when the brain/CNS is being sensitively driven by the

body/environment in a specific way, while still maintaining that the involvement of the body/environment is here merely causal rather than genuinely constitutive. (Clark 2009 p. 987)

The following picture emerges from all this. On the one hand, consider non-conscious, non-occurrent dispositional states like beliefs; according to EM, they can be extended so that their material basis involves, as constituents, parts of the world outside an agent's body. Contrast these with conscious events like experiences. There is no parallel argument here to the argument supporting EM; the contribution of the environment to conscious events is causal, not constitutive.

3. Tetris, blocks copying task – cases of ECM?

As mentioned above, Clark regards the original argument as relying partly on some key thought-experiments. I have already described one of these, about Otto and his notebook. Now it's time to mention the other, which opens the 1998 paper. Imagine three scenarios: (1) an agent playing Tetris mentally imagines rotating the blocks (2) a person playing Tetris pushes a button to rotate the blocks on the screen that she is watching (3) an agent in the cyberpunk future has a neural implant which can perform the rotation as fast as the computer screen in the previous case, and she can choose between either using the implant, or using mental imagery to rotate the blocks.

Clark and Chalmers note that the three scenarios seem to be analogous. If we are inclined to classify (3) as an example of a cognitive process, then we should do the same with (2), which has a computationally analogous structure “distributed across agent and computer instead of internalized within the agent” (Clark and Chalmers 1998 p. 7). The setup of the whole thought experiment is analogous to the Otto/Inga/Martian thought experiment. The fact that Otto's and Inga's situations are on a par is supported further by the intermediary example of the Martian with an internal bitmap storage. Similarly, the fact that the internal mental rotation of the blocks is analogous to the rotation done by agent-plus-screen is supported by the intermediary example of the cyberpunk agent with the implant. The Parity Principle applies in the same way. Indeed, the Parity Principle itself (though not under this name) was formulated in the original paper as the lesson of the Tetris case. The idea is that the rotation of the blocks on the screen accompanied by the agent's visual experience functions in the same way as the implant, and as the internal mental rotation.

On the basis of this, the second scenario of the Tetris case, the agent watching the screen while she is rotating the blocks, looks like a candidate for an extended mental feature which is just as good as the Otto case. Now notice that the Tetris case does not concern a non-conscious dispositional state: it concerns an event that is part of the stream of consciousness, namely the

visual experience of watching the screen. It may be countered that the event is not essentially conscious – it involves the kind of functional process that say, a machine could exemplify. I will address this point later (in section 5), but for now I note that there is no doubt that *in us*, this event is conscious. Maybe machines could do analogous things, but Clark and Chalmers's example concerns a human agent watching a screen, and in humans, perceptual experiences of this kind are part of the stream of consciousness. And if the proper assessment of the case requires us to say that the cognitive process is distributed over agent and screen, then it seems that here we have an *extended conscious event*. Contrary to what Clark says, the argument for EM seems to function also as an argument for ECM.

Clearly, this is still too quick. Given that Clark is resolutely sceptical about ECM, he presumably doesn't see the Tetris case as supporting the extension of consciousness. Starting in the next section, I will consider some possible explanations for this. But before I move to these explanations, I would like to mention another candidate case for supporting ECM which Clark cites in several writings. It is an experiment conducted by Dana Ballard and coworkers (1997) which involves the task of copying a pattern of coloured blocks. Ballard et al. found that subjects memorised only either the position or the colour of a particular block which was being copied, not both. When the other piece of information was needed, subjects glanced back at the model.

The moral, according to Clark, is that the movement of the eye, and the blocks themselves, play an identifiable functional role in solving a task. The blocks themselves store the information about their position and colour by simply having these features, and this is on a par with an internal representation of position and colour. Further, “changing gaze is analogous to changing the memory reference in a silicon computer.” (Ballard et al. 1997. p. 725, quoted by Clark 2008b, p. 45). Clark doesn't explicitly evoke the Parity Principle, but it seems that it easily applies: if the storing and recall of the information from internal memory are part of a cognitive process, then the functionally equivalent information-stored-by-the block and changing gaze are also part of a cognitive process.

Clark thinks that the Ballard study and the Otto story support the same thesis of “extended functionalism”.

The [Otto] argument ... echoes Ballard et al.'s in depicting a larger integrated system as the extended machinery whose computationally salient states, properties and transformations are supposed to explain specific problem-solving performances. (Clark 2008a, p. 46)

Notice that again, we have here an occurrent conscious process of problem-solving, rather than a non-conscious dispositional state. Just as before, when I discussed the visual experience of

watching the Tetris screen, it could be countered that this process is not essentially conscious, since the same task could be carried out by a zombie or by a machine. As before, I promise that I will address this point later, and note for the time being only that in the actual subjects of experiment, the process is clearly conscious. So the actual case is an example of a conscious event.² Yet it seems that the process extends beyond neural resources: it contains the movement of the eye and the blocks that store the information. Do we then have another example of an extended conscious event, a support for ECM?

4. Occurrent but non-conscious processes?

We have seen that Clark treats both the Tetris case and the Ballard et al. case as analogous to the Otto case in important respects: in the application of the Parity principle, or as examples of extended functionalism – which both arguably express the same idea. Since actual examples in both the Tetris case and the Ballard et al. case involve occurrent conscious processes of human cognizers, this suggests, at first sight, that we have an argument for Extended Conscious Mind that is exactly analogous to the argument for Extended Mind. In what follows, I will look at possible explanations of why, despite this appearance, Clark resists the idea of ECM.

If we want to resist ECM, we have at least two options. First, we could deny that the Tetris and the Ballard et al. cases provide bona fide cases of extension; this will be considered in section 6. Alternatively, if the Tetris case and the Ballard et al. case are indeed analogous to the Otto case and provide bona fide cases of extension, then we need an explanation of why they are not extensions of *consciousness*.

In Clark's paper "Spreading the Joy" (Clark 2009), which addresses the issue of extended consciousness at length, the examples of extended non-conscious mental features all concern non-occurrent states like beliefs:

the Clark/Chalmers case for EM (...) targets only non-conscious mental states and processes, where long-term informational poise—rather than online informational access and integration—seems to be what counts. (Clark 2009, p. 983)

In contrast, the Tetris case and the Ballard et al. case involve occurrent processes of problem solving (where online informational access and integration do seem to matter a great deal). Of course, not all occurrent processes are conscious. But we haven't yet seen an explanation of how occurrent processes can be non-conscious, or if they are conscious, why their extension does not

² Thanks to the editors for pressing me on this point.

involve an extension of consciousness.

As a starting point, consider the following brief remark by Chalmers in the foreword to Clark's *Supersizing the mind*:

But then, what about the big question: extended consciousness? The dispositional beliefs, cognitive processes, perceptual mechanisms, and moods considered above all extend beyond the borders of consciousness, and it is plausible that it is precisely the non-conscious part of them that is extended. (Chalmers 2008, p. xiv)

Cognitive processes and perceptual mechanisms are occurrent processes, yet, according to this remark, they have a part that is non-conscious and can be extended. Could the non-conscious part refer to the early processing stages of perceptual mechanisms which, on some suggestions, precede what is hypothesized to be the emergence of consciousness? On a closer look, this seems unlikely. There is no suggestion at all in the descriptions of the Tetris case that the extension attaches specifically to pre-conscious processing. We had better look for another sense in which processes extend beyond the borders of consciousness.

5. The separation of consciousness from the functional/representational

Here is another interpretation of the claim that perceptual mechanisms stretch “beyond the borders of consciousness”: it reflects the view that the functional/intentional aspect of mental features can be entirely divorced from their conscious character. We can disregard the conscious features of, say, a visual experience, and focus exclusively on its functional characterisation. The functional part is “beyond the borders of consciousness”, since it can be considered independently of any conscious features.

Chalmers says that mental extension plausibly concerns the non-conscious part, and this is in harmony with the conscious/functional separability interpretation. The case for EM is based on the Parity Principle or some version of “extended functionalism”, so the relevant occurrent processes are indeed characterised in functional terms. We have already encountered the idea above that the perceptual process in the Tetris case, and the problem-solving process in the Ballard et al. case, could be exemplified by a functionally equivalent but non-conscious zombie or machine. So the process that gets extended is not conscious, but a functionally characterised process.

I think that a belief in the separability of the conscious and the functional is the most likely explanation of Clark and Chalmers's resistance to extending consciousness. But there are two problems with this idea. First, the conception itself is rather problematic for independent reasons.

Take for example the visual experience in the Tetris case. Having the experience enables the agent to solve the task at hand (finding out whether the blocks fit into the slots) because it represents the moving shapes. So the representational features of the experience are intimately connected to its functional role in the task. The separability conception asks us to separate the conscious character from the functional/representational features. What it's like to have the experience will be independent from, say, its looking to the agent that a long block is above a narrow vertical slot. Many would find this implausible. Those who do, can sign up to theories of perception which oppose the idea that the conscious character of perceptual experiences is separable from their representational (or presentational) features: representationalists, naive realists, defenders of the phenomenal intentionality view. I propose that defenders of all these views would object to the idea that we can talk about the nature of a perceptual experience while entirely bracketing its conscious character.

There is another, dialectically more pressing reason why the separability conception fits rather uneasily with the overall picture presented by Clark. If we revisit the arguments of the 2009 paper against the Extended Conscious Mind, it seems that they will disqualify the Tetris and the Ballard et al. cases from extension even if we assume that the conscious and the functional are separable.

Recall what Clark describes as an example of the “standard view”: the view that motion perception is realised in a certain part of the brain (and possibly the central nervous system). Defenders of ECM attack this idea, but Clark claims that none of the arguments for ECM work. So until further notice, I will assume that the standard view remains intact. Note that this argument talks about the visual experience as a whole. I suppose it's possible that there is a neural correlate of conscious character that is actually separable from the neural correlates of the representational and functional features of an experience. But there is absolutely no suggestion in the evidence considered for the standard view that only the neural correlate of the conscious character is localised, and the neural correlates of representational and functional aspects are not. On the contrary. The evidence Clark cites in defense of the standard view concerns selective deficiency in visual tasks, observed both in a human subject and in macaque monkeys whose Middle Temporal area was removed (Marcar et al. 1997). The tasks identify visual experiences in terms of their functional properties, not in terms of an elusive conscious character. What the macaque monkeys are provably missing is not a conscious quality, but the ability to function in a certain way. If the standard view is safe from ECM, then visual experiences as a whole (understood as conscious events playing a certain functional role eg. in visual tasks) are localised in the brain.

Apply this to the Tetris case which actually involves the same type of experience: motion

perception. On the standard view, this experience – which enables the subject to solve certain tasks – is realised in a certain part of the brain. The standard view is not threatened by arguments for ECM, according to Clark, so I will assume it's correct. It looks like we should say that the perceptual experience of the Tetris screen is firmly in the head.

Consider next the argument from dynamic entanglement. Clark thinks that intuitively, “visual experience may at each moment depend solely on the complex brain activity caused (instrumentally) by the ongoing engagements with body and world.” (Clark 2009, p. 977) Again, this remark treats the experience as a whole, and there is no reason why this wouldn't apply to the Tetris case. Consider the further proposed “twin test”: would a neural duplicate of the Tetris player (without the screen) share his experience? If there is no argument to the contrary (and Clark seems to think that there isn't), then the duplicate will have a perfect hallucination of the screen. But the perfect hallucination will be sufficient to solve the task, that is, to answer the question of whether a block fits into a pattern or not – it will simply act as an unusually vivid and fast manipulable mental image. In other words, we have an argument here that the visual experience, identified partly in terms of its task-solving potential, has its material basis in the brain/CNS. The screen is of course important in sustaining the experience, but its contribution is best viewed as “merely causal rather than genuinely constitutive”. (Clark 2009 p. 987)

It seems that Clark's arguments against ECM locate in the head the material basis of the entire perceptual experience, not only the neural correlates of a potentially separable conscious character. So the option that the Tetris and the Ballard et al. cases are extensions but not really extensions of consciousness is undermined by these arguments. The separability thesis does not explain why the Tetris and the Ballard et al. cases are not cases of extended consciousness. This directs our attention to the other possible strategy involving the Tetris and the Ballard et al. cases: deny that they are bona fide cases of extension. I will consider this option in the next section.

6. The Tetris and Ballard et al. cases are not extended

In this section, I will try to find differences between the Tetris and the Otto cases that could suggest that only the second is a genuine case of the Extended Mind. First, a point mentioned by Clark in *Supersizing the Mind*: he notes that the Tetris player can rely on the screen only when the computer is available, and such limited availability is not true for our usual cognitive resources. The Otto example takes the argument “one step further” (Clark 2008b, p. 78), Clark claims, because it addresses the issue of portability. Now this is indeed an important difference between the cases as presented, but it hardly goes to the heart of the matter. There is no in principle obstacle to designing a scenario where the external resource employed in a conscious process is just as portable and

available as Otto's notebook. Issues of portability will not disqualify consciousness from extension.

Here is a second attempt. Clark and Chalmers state that cases like Tetris and others support an extension of “cognitive processing” into the environment. However, they think this falls short of the extension of the “mind”:

Perhaps some processing takes place in the environment, but what of mind? Everything we have said so far is compatible with the view that truly mental states – experiences, beliefs, desires, emotions, and so on – are all determined by states of the brain. Perhaps what is truly mental is internal, after all? (Clark and Chalmers 1998, p.12)

This is followed by the Otto case which, according to the authors, takes things a step further: not just processing, but the mind itself is extended. The distinction is standardly taken up in the literature. For example, in a survey on “Varieties of Externalism”, Adam Carter, Jesper Kallestrup, Orestis Palermos and Duncan Pritchard (2014) distinguish between the hypothesis of *extended cognitive processing* and the thesis of *extended mind*. So one explanation of why Clark doesn't regard the Tetris case as supporting ECM is that he thinks it supports merely the extension of *cognitive processing*, but not the extension of the *mind* (whether conscious or not).

But on closer examination, the difference between the two theses becomes elusive. For what is “cognitive processing”, if not a process that involves a succession and development of mental (cognitive) events? And how could the process be extended without its constituent events being extended? Certainly, there is a difference between the kinds of mental features considered in the Tetris and the Otto case: in one, we have a conscious occurrent process of problem-solving, in the other, we have a non-conscious dispositional state. But what is *claimed* about the mental features is the same: that they are distributed over the agent and an external factor. Carter et al. suggest that the arguments for the two theses are also the same: a kind of common-sense functionalism embodied in the Parity Principle. The best Carter et al. can do to state the difference between the theses is to call attention to their different *focus*: processes or states. But that is just a difference between the kinds of mental features that are claimed to be extended, and not the nature of extension. Saying that the Tetris case supports the extension of “cognitive processing” does not, in itself, explain why it doesn't support ECM.

So issues of portability, and the contrast between processes and states will not do the trick. Here is a third attempt: the difference between the Tetris and the Otto case concerns not the kind of feature that is extended, but the nature of the extension. It is only in the Otto case that the external factor is a genuine *constituent* of the material basis of a mental feature. In contrast, in the processing

cases, the extension falls short of being a constituent, and plays only some sort of intimate causal role. An example of this kind of view can be found for example in Margaret Wilson's widely cited paper on varieties of situated cognition. Wilson characterises Kirsh and Maglio's (1994) work on epistemic action performed in playing Tetris (which was Clark and Chalmers's inspiration for the Tetris example) and the Ballard et al. 1997 experiment as examples of off-loading cognitive work to the environment. In these cases, the environment plays a causal role, but it is not actually part of the of the cognitive process.

This, I believe, is the most plausible option for the defenders of EM to take. However, this position requires a thorough revision of Clark's description of the Tetris and the Ballard et al. cases, and of the application of the Parity Principle. We better not say that the process in the second scenario of the Tetris case is “distributed” over agent and external factor. And we should either give up the idea that the Parity Principle is suitable for liberating us of bio-chauvinistic prejudices and for identifying extended features, or rethink carefully how it exactly applies to the Tetris and the Ballard et al. cases. For example, if the visual experience of the screen *alone* (though caused by the presence of the screen) is functionally on a par with the mental rotation of images, then the Tetris case does not offer a *prima facie* case for extension.

7. Reformulating the debate about EM

I have argued that there is a tension between the claim that EM does not apply to conscious states, and the treatment of the Tetris and the Ballard et al. cases found in Clark's writings. I think that the best resolution of the tension is to state clearly that the Tetris and the Ballard et al. cases do not support EM, even if Otto's case does. In this closing section, I would like to propose a reformulation of the debate that explains why we should focus on the Otto case rather than on the other cases when debating the issue of EM.

The extended mind thesis is sometimes put by saying that the mind “literally extends beyond the body”. Many people find this less than helpful. If the mind literally extends, then presumably it has a certain weight, a shape, possibly has holes in it, could or could not fit into a small cabin bag, and so on. None of these sound like features of a mind. The situation somewhat improves if we talk about the minimal material basis of the mind. But there is another way of phrasing the debate which steers clear of the issue of extension and location. It is simply this: given his relationship to his notebook, *does or does not Otto have the belief that MoMA is on 53d street?* Defenders of the Extended Mind thesis will say that Otto does have the belief. Opponents will deny this. We are not hereby talking about the location or extension of “the mind”, nor do we need to assume that the mind is “a thing” – we simply ask of a cognitive agent, in a given situation, whether he possesses a

certain mental feature or not. We can also put aside all concerns about derived and non-derived representations: the attribution of the belief is to a subject who undoubtedly has cognitive states.

This is a question with serious implications to our understanding of the nature of standing states. The idea that standing states like beliefs can be characterised in functional terms, or their nature is given by a set of dispositions, is very attractive. But it can be argued that what matters is the macro-functional, rather than the micro-functional role of these states. This can be illustrated by Otto's case as follows. The information in Otto's notebook will function *by-and-large* the same way as Inga's biological memory, insofar as it entails the same action when combined with the same desire. But there are several details that are different: for example, Otto recalls information through the mediation of a perceptual experience, while this isn't true of Inga. If Inga abandons a belief, many of her other beliefs will be automatically updated; while if Otto deletes an entry in his notebook, he has to make sure separately that other entries are synchronised. The idea is, however, that the nature of a standing state is determined by the kind of coarse-grained equivalence that does exist between Inga and Otto, and the kind of fine-grade differences that also exist do not matter (for further arguments, see Gertler 2007 and Sprevak 2009).

If we agree that Otto has the relevant belief, it is plausible to speculate that the material basis of the belief includes the notebook. But note that this is not essential for setting up the debate. We could equally ask whether Otto's brain-in-a-vat counterpart has the belief. The crucial question is whether a somewhat non-standard dispositional profile is enough to attribute a standing state. Similar questions will arise about other standing states. Take a person who always relies on her notebook for keeping appointments – does she or does she not have the intention to attend a certain meeting?

Notice that there is no similar issue for the Tetris and the Ballard et al. cases. We are not asking the following question: given its relationship to the screen, does or does not the Tetris player have the visual experience of the screen? We all agree that she does. In the putative cases of extension, defenders and opponents of extension will attribute the same mental features to the agents.

I argued elsewhere that there is a fairly innocuous sense of “extending” conscious states that is arguably *not* at issue in the EM debate (Farkas 2012). Imagine that silicon wiring which runs outside the skull can provide a perfect replacement for some neurons that are plausibly part of the material base of some conscious event. The replacements are so good that subjects with the wires installed don't notice any difference in their mental life. There is no suggestion that a merely silicon-based creature would also be conscious; only a few neurons are replaced. I proposed that the possibility of such a perfect “prosthesis” is compatible with a whole range of philosophical views

of the mind, including those of Clark and of many of his opponents (whether it's empirically possible is another question). So what gets (most) people going in the debate about EM is not extension as mere prosthesis, but something else. I suggested that one important issue is precisely the one presented above: whether standing states can be attributed on the basis of a non-standard dispositional profile. This is a question that arises for the Otto case, but doesn't arise for the Tetris and the Ballard et al. case, as they are normally presented. Hence we will do well if we focus on the first case when we try to decide to possibility of extending the mind.

References

- Ballard, Dana H.; Hayhoe, Mary M.; Pook, Polly K. & Rao, Rajesh P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences* 20 (4):723-742.
- Carter, J. Adam ; Kallestrup, Jesper ; Palermos, S. Orestis & Pritchard, Duncan (2014). Varieties of externalism. *Philosophical Issues* 24 (1):63-109.
- Chalmers, David J. (2008). Foreword to Andy Clark's *Supersizing the Mind*. In Andy Clark: *Supersizing the Mind*. Oxford University Press: ix-xvi
- Clark, Andy (2005). Intrinsic content, active memory, and the extended mind. *Analysis* 65 (285):1-11.
- Clark, Andy (2008a). Pressing the flesh: A tension in the study of the embodied, embedded mind? *Philosophy and Phenomenological Research* 76 (1):37-59.
- Clark, Andy (2008b). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Clark, Andy (2009). Spreading the joy? Why the machinery of consciousness is (probably) still in the head. *Mind* 118 (472):963-993.
- Clark, Andy & Chalmers, David J. (1998). The extended mind. *Analysis* 58 (1):7-19.
- Farkas, Katalin (2012). Two Versions of the Extended Mind Thesis. *Philosophia* 40 (3):435-447.
- Gertler, Brie (2007). Overextending the mind? In Brie Gertler & Lawrence Shapiro (eds.), *Arguing About the Mind*. Routledge. pp. 192--206.
- Kirsh, David & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science* 18 (4):513-49.
- Marcar, V., J. Zihl, and A. Cowey 1997: 'Comparing the Visual Deficits of a Motion Blind Patient with the Visual Deficits of Monkeys with Area MT Removed'. *Neuropsychologia*, 35, pp. 1459-65.
- Sprevak, Mark (2009). Extended Cognition and Functionalism. *Journal of Philosophy* 106 (9):503-527.
- Wilson, Margaret (2002). Six Views of Embodied Cognition. *Psychonomic Bulletin and Review* 9 (4):625--636.