

Draft of chapter to appear in in Gardini, N., Jacobs, A. X., Morgan, B., Omri, M.-S., and Reynolds, M. (eds.) *Minding Borders*, Legenda, Oxford.

**The Edge of Thought:
Extended Cognition and the Border between Mind and World**

**Michael Wheeler
University of Stirling**

1. Minds in Space

Where is the mind? If we understand this question literally – that is, as requiring something like spatial coordinates as an answer – some people will complain that it's simply preposterous. Psychological phenomena, such as thoughts, chains of reasoning, emotions, and feelings, are just not the right sorts of 'things' (they will say) to have such coordinates. Indeed, when Robert Flack and Donny Hathaway posed the immortal question 'Where is the love?', they didn't expect an answer of the form 'two metres to the left of the door, just above the sofa, with a volume about the size of a football'. And it's not only great soul singers who will be queueing up to pour scorn on the idea that mind literally has spatiality. The frowning line of naysayers will also include a motley assortment of philosophical heavyweights, including the much-maligned substance dualists, who hold that mind is a non-physical entity which is distinct from the material world and so certainly not extended in space, and a range of thinkers who, despite their differences, would all be moved to protest, perhaps with a Heideggerian, Rylean, or later Wittgensteinian swagger, that even though substance dualism is a confused doctrine, to think of a mind as a spatially locatable entity is to commit a kind of logical error in one's use of the very concept of mind.

All that said, and with apologies to Roberta, Donny and their philosophical friends, I think our opening question is not only perfectly reasonable, but has a surprising and controversial answer. Here's a quick and dirty argument for why it is reasonable. Substance dualism is officially unpopular with most philosophers and scientists, and a natural (and naturalistic) alternative is to claim that minds are (in some way) realized in the physical world, rather than in some non-physical substance. But if minds are (in some way) realized in the physical world, then (in some sense) they must take up space. And if minds (in some sense) take up space, then (in some sense) it must be reasonable to ask where in space they are. QED. Well, not quite. My irritating use of qualifying parenthetical phrases in the foregoing argument emphasizes just how under-specified that argument is. Nevertheless, the key point is, I think, secure enough: if we are interested in the realizers of psychological states

and processes, and if those realizers are physical elements on a metaphysical par with doors, sofas and footballs (if one dislikes the ‘thing-y-ness’ of these examples, one can simply switch to examples of physical systems, networks, or even flows), then it is perfectly reasonable to wonder where in space those realizers are, just as it is perfectly reasonable to ask about the spatial position and extent of more pedestrian physical phenomena. In other words, if the question ‘where is the mind?’ is interpreted as meaning ‘where are the physical elements that realize psychological states and processes?’, then it isn’t preposterous at all. Indeed, one could press the point about realizers while also thinking that, in some psychological contexts (such as that of revealing the intentions of a lover whose commitment is in doubt), it would be an inappropriate way of addressing a ‘where’ question. Nevertheless, once the focus *is* on realizers, it seems that we can meaningfully talk about minds that take up space.

One consequence of this state of affairs is that it also makes sense to conceive of any particular mind as having some sort of spatially locatable outer limit that marks the border between mind and world – the edge of thought, as one might put it. But, of course, given that such a border exists, there is, in principle anyway, room for genuine disagreement over just where in space it is. That’s the ‘where’ issue that this paper addresses.¹

2. Extended Cognition

It is time for us to re-ask our opening question, but this time without all the defensiveness. Here we go. Where is the mind? More specifically, where is the edge of thought, the border between mind and world? One answer with plenty of

¹ In my opening paragraph, I suggested that thinkers such as Heidegger, Ryle and the later Wittgenstein were representatives of a nonspatial view of mind. But now that we have separated out questions about physical realizers from other questions in the vicinity of cognition, it might seem that although these thinkers were themselves focussed on a level of analysis separate from realization, a level at which thinking of psychology in terms of spatiality is, as they might argue, misguided, nevertheless their philosophical frameworks might, without violence, be mined for insights about, or for guidance regarding, questions about realizers. And, in a sense, that’s precisely what I shall do later in this paper, by calling on Heidegger. However, as a matter of good methodology, it seems abundantly clear that one should begin by being suspicious of any analysis that appeared simply to *read off* a view about realizers from a view pitched at an altogether different level of analysis. For more on this sort of issue, see e.g. section 4 of Michael Wheeler, ‘The Revolution will not be Optimised: Radical Enactivism, Extended Functionalism and the Extensive Mind’, *Topoi* (2015), SpringerLink Online First, DOI: 10.1007/s11245-015-9356-x.

contemporary momentum (see below) is that the mind-world border is stationed at the sensory-motor interface between the organic body and the beyond-the-body environment. For present purposes, I am not going to fuss over the genuine distinction that exists between saying that the mind is located in (or perhaps is co-located with) the brain or central nervous system, and saying only that the mind is located inside the skin. Whether we go for skull or skin as our mind-world border, what we confront is a version of a view that some philosophers call *vehicle internalism*. To explain: in the present context, the term ‘vehicle’ is practically equivalent to the term ‘realizer’. According to vehicle internalism, then, psychological states and processes are realized by physical machinery that is always inside the skull or skin. I would fearlessly risk my treasured collection of early punk rock singles on the bet that, right now, vehicle internalism is overwhelmingly the default view in cognitive neuroscience, cognitive psychology, developmental psychology, cognitive-science-friendly philosophy, and indeed in many further disciplines and communities (including some situated in the wider humanities). Indeed, given all those ‘pictures of the brain thinking’ that have been produced over the past few years by modern neuroimaging techniques, one might reasonably wonder where else the material machinery of mind could possibly be. So any answer to our ‘where’ question that departs from vehicle internalism surely counts as surprising and controversial, although of course it won’t surprise, or attract passionate disagreement from, everyone.

Just such an answer has been defended in and around cognitive science recently by advocates of the so-called *hypothesis of extended cognition* (henceforth *ExC*).² Rejecting vehicle internalism, fans of *ExC* endorse the competing position of vehicle *externalism*. In other words, according to *ExC*, the physical machinery of mind *sometimes* extends beyond the skull and skin. More precisely, if *ExC* is true, then there are actual (in this world) cases of intelligent thought and action, in which the material vehicles that realize the thinking and thoughts concerned are spatially distributed over brain, body and world, in such a way that certain external (beyond-the-skull-and-skin) factors are rightly accorded fundamentally the same cognitive status as would ordinarily be accorded to a subset of your neurons. Eye-catching

² The original presentation and case for *ExC* was in Andy Clark and David Chalmers, ‘The Extended Mind’, *Analysis*, 58: 1 (1998), 7-19; see also Andy Clark, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* (New York: Oxford University Press, 2008). For a more recent collection that places the original Clark and Chalmers paper alongside a range of developments, criticisms and defences of *ExC*, see Richard Menary, ed., *The Extended Mind* (Cambridge, Mass.: MIT Press, 2010). Throughout this chapter, I will use the terms ‘mind’ and ‘cognition’ interchangeably. Although this irritates some people, it strikes me as standard practice in cognitive science. From this perspective, ‘extended cognition’ and ‘the extended mind’ are alternative names for the same view.

examples of external elements that advocates of ExC often take to have such cognitive status include smartphones, tablets and at least some instances of wearable computing, but, in the end, nothing much hangs on such feats of contemporary technological wizardry. Less fancy items such as notebooks (the old-fashioned kind), tally sticks and abacuses would, under the right circumstances, do just as well.

Of course, on all (sane) views, there is a sense in which minds reach beyond the skin, in that they secure epistemic contact with what our thoughts are (typically) about, namely the external world. As Heidegger puts it:

We do not represent distant things merely in our mind... so that only mental representations of distant things run through our minds and heads as substitutes for the things. If all of us now think, from where we are right here, of the old bridge in Heidelberg, this thinking toward that location is not a mere experience inside the persons present here; rather, it belongs to the nature of our thinking of that bridge that in itself thinking gets through, persists through, the distance to that location. From this spot right here, we are there at the bridge – we are by no means at some representational content in our consciousness.³

One could mistake turns of phrase such as ‘thinking... persists through... the distance to [a remote] location’ as endorsements of ExC, but, of course, jettisoning the idea that thinking cannot reach beyond a veil of inner representations (which is what Heidegger is up to here) is not, in and of itself anyway, sufficient to establish ExC, since ExC concerns not the correct characterization of the intentional or epistemic reach of thought, but rather the location of its underpinning physical machinery.⁴

³ Martin Heidegger, ‘Building, Dwelling, Thinking’, in Martin Heidegger, *Poetry, Language, Thought*, trans. by Albert Hofstadter (New York: Harper and Row, 1971), pp.145-61. Reprinted in Martin Heidegger, *Basic Writings*, ed. by David Farrell Krell (London: Routledge, 1978), pp.343-63. Quotation from p.358 of the Krell edition.

⁴ I say ‘in and of itself’ here because I don’t rule out the possibility that additional arguments might take us from an analysis of the structure of, say, intentionality to ExC. For arguments in this area – minus the antirepresentationalism – see Mark Rowlands, *The New Science of the Mind: From Extended Mind to Embodied Phenomenology* (Cambridge, Mass.: MIT Press, 2010). For the view that rejecting any sort of robust representational account of basic cognition is a key step in securing something like ExC, see: Daniel D. Hutto and Erik Myin, *Radicalizing Enactivism: Basic Minds without Content* (Cambridge, Mass.: MIT Press, 2010); Daniel D. Hutto, Michael D. Kirchhoff and Erik Myin ‘Extensive Enactivism: why Keep it all in?’, *Frontiers in Human Neuroscience* 8: 706 (2014), 1-11. For a critical response to the latter, antirepresentational view, see Wheeler, ‘The Revolution will not be Optimised’.

To highlight more clearly what one ought to mean when one claims that the machinery of mind is extended, consider the way in which a skilled bartender may achieve the successful delivery of a large and complex order of drinks.⁵ Fulfilling such an order can strike the casual observer as a relatively daunting task, especially if the memory-relevant resources available to the bartender are thought to be restricted to inner storage and recall. However, it is a fortuitous fact that different kinds of drink often come in differently shaped glasses. (Think of cocktails.) The cognitive ecology of the bar is thus characterized by some relatively persistent physical structures (the differently shaped glasses), plus some culturally established norms (the specificity of kind of drink to shape of glass). What novice bartenders learn to do, under the tutelage of their experienced colleagues, is to retrieve the correct glass for each drink as it is requested, and to arrange the differently shaped glasses in a spatial sequence that tracks the temporal sequence of the drinks order. From the perspective of the purely organic cognitive resources available, what the bartender has learnt to do, in replicating this culturally transmitted practice, is to exploit her physical environment in order to outsource complexity and so reduce the burden on inner processing. In effect, she transforms what might have been a highly challenging memory task into a simpler (roughly) perception and association task.

This is a compelling example of what Andy Clark has dubbed *cognitive niche construction*, the process by which human beings build external structures that, in combination with culturally transmitted practices, transform problem spaces in ways that promote (or sometimes obstruct) thinking and reasoning.⁶ As such, it is at least an example of what is now commonly called *embedded cognition* – cognition that itself remains wholly internal, even though it exhibits subtle and complex patterns of causal dependence on the beyond-the-skin environment. The advocate of ExC wants to go further than this, however, by demanding that it's not merely that the bartender's successful delivery of the drinks order depends *causally* on her non-neural bodily movements (the manipulation of different glasses so as to produce an order-tracking spatial arrangement), as well as on material environmental structures (the actual spatial arrangements of different glasses). Rather, those non-neural bodily movements and material environmental structures are themselves revealed to be bona fide *constituents* of the realizing physical machinery of the cognitive trait of memory. If this second interpretation is correct, our distributed embodied-

⁵ See King Beach, 'The Role of External Mnemonic Symbols in Acquiring an Occupation'. In *Practical Aspects of Memory*, ed. by M. M. Gruneberg and R. N. Sykes (New York: Wiley, 1988), pp.342-346.

⁶ Beach's analysis of bartender behaviour is developed as an example of cognitive niche construction by Clark in his *Supersizing the Mind*; see also Michael Wheeler and Andy Clark, 'Culture, Embodiment and Genes: Unravelling the Triple Helix', *Philosophical Transactions of the Royal Society Series B* 363 (2008), pp.3563–3575.

bartender-plus-glasses system is a paradigmatic case not of embedded, but of extended, cognition. The border between mind and world has been relocated to a point beyond the skin.⁷

3. Perception, Action, Mind, World

In this chapter I am not going to present arguments in direct support of ExC.⁸ Instead, I am going to describe a seemingly powerful objection to the extended view and then show how responding to that objection eventually leads us to say something that is (with respect to the orthodox view in and around cognitive science anyway) profoundly unexpected about the border between mind and world. The worry in question is nicely expressed by David Chalmers (ironically, of course, one of the original architects of ExC) in the following short passage. 'It is natural to hold that perception is the interface where the world affects the mind, and that action is the interface where the mind affects the world. If so, it is tempting to hold that what precedes perception and what follows action is not truly mental.'⁹ Here is the picture that Chalmers describes as 'natural' and 'tempting'. Where perception is the focus, there are events and states of affairs that occur in the world, and these (conceptually, and perhaps typically temporally too) come before the experiences of, and the beliefs about, those worldly states of affairs that occur in the mind, and that are produced via perception. Perception is thus the world-to-mind channel by which certain beyond-the-mind happenings generate certain psychological happenings. Where action is the focus, there are desires and related thoughts about how to satisfy those desires that occur in the mind, and these (conceptually, and perhaps typically temporally too) come before the changes to the states of affairs in the world that occur so as to satisfy those desires, and that are produced via action. Action is thus the channel by which certain psychological happenings generate certain beyond-the-

⁷ My chosen way of expressing the difference between embedded cognition and ExC, that is, in terms of a causal-constitutive distinction, is due originally to Fred Adams and Kenneth Aizawa, *The Bounds of Cognition* (Malden, MA and Oxford: Blackwell, 2008), although it is now widespread in the literature.

⁸ The most-discussed argument for ExC, which is based on what is called the parity principle, is to be found in the original Clark and Chalmers treatment of the issue. For alternative pro-ExC arguments that claim various advantages over parity-based considerations, see e.g.: Richard Menary, *Cognitive Integration: Mind and Cognition Unbounded* (Basingstoke: Palgrave Macmillan, 2007); John Sutton, 'Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process' in Richard Menary, *The Extended Mind*, pp.189-225; Michael Wheeler, 'Revolution, Reform, or Business as Usual? The Future Prospects for Embodied Cognition', in *The Routledge Handbook of Embodied Cognition* ed. by Lawrence Shapiro (Abingdon and New York: Routledge, 2014), pp. 374-83.

⁹ David Chalmers, Preface to Andy Clark, *Supersizing the Mind*, p.xi.

mind happenings. Put all this together, and that's why the border between mind and world is located at the perception-action interface.

This is indeed an intuitively compelling picture, but how, exactly, does it generate an objection to ExC? Consider, once again, our complexity-outsourcing bartender. An attentive observer might be moved to make two observations about the bartender's order-delivering activities: (i) prior to the empty glasses being placed physically in their appropriate spatial positions, the bartender plausibly has various order-related desires, and thoughts about how to satisfy those desires, that kick-start and organize her bodily movements; (ii) prior to making the right drinks, the bartender must look at the glasses before associating their particular shapes with specific kinds of drinks. The first of these observations provides an example of Chalmers' mind-to-world action channel, with bodily movement understood as the border crossing between mind and world, while the second provides an example of Chalmers' world-to-mind perception channel, with visual sensing understood as the border crossing between world and mind. But if this analysis is correct, then the relevant beyond-the-skin elements, namely the glasses in their spatial arrangements, are on the *wrong side* of the perception-action interface to count as realizing vehicles of the bartender's cognitive states and processing. Those external elements scaffold, but they do not realize, bartender-cognition. In other words, our brain-body-environment drinks-order system is an example of an embedded, rather than an extended, mind.

How should the fan of ExC respond? In their classic treatment, Clark and Chalmers¹⁰ argue that the correct reply to the foregoing sort of worry is to claim that although our intuitive inclinations are to categorize an organic agent's use of her eyes as perception and her intentional deployment of bodily movements as action, those inclinations are not always reliable. Indeed, to assume that they are always reliable would be to beg the question against ExC. To explain: in effect, Clark and Chalmers suggest that although the perception-action interface may constitute the border between mind and world, that interface is not always co-located with the organic sensorimotor boundary. Thus, and to apply this suggestion to our bartender example, from an extended mind perspective, the bartender-plus-glasses assemblage constitutes a cognitive system in its own right, so the salient events that involve the bartender first moving the glasses into sequence and then looking at the sequentially ordered glasses all happen *inside* that cognitive system. But if that is right, and if the perception-action interface fixes the border between mind and world, then the organic sensorimotor boundary cannot be co-located with the perception-action interface. The latter has to be at the edge of the cognitive system, not inside it. So how are we to think about the bartender's deployment of bodily manipulations to move the glasses around and her use of her eyes to access information about the glasses? The idea is that, in comparison with ordinary cases of inner cognition, we

¹⁰ Clark and Chalmers, 'The Extended Mind'

should think of these occurrences as more akin to wholly internal information processing – as cases of internal memory recall and introspection perhaps – than to the world-to-mind and mind-to-world events that are indicative of perception and action. In response to the obvious complaint that the bartender’s access to the glasses involves a distinctive phenomenal experience that is not present in the case of wholly inner processing, Clark and Chalmers appeal to the science fiction example of Arnold Schwarzenegger’s Terminator character, as depicted in the James Cameron film. Information that the Terminator retrieves from his memory is presented to him in his conscious visual field. The fact that this access system presumably has an accompanying visual phenomenology does not seem to prevent it from realizing the Terminator’s memory, so why should it have that sort of consequence in the case of our distributed bartender-plus-glasses system? The answer, Clark and Chalmers suggest, is that it shouldn’t.

It is hardly news to note that the defence of ExC just canvassed involves a good deal of revision of our natural way of thinking about the border between mind and world. Of course, we should not be afraid of such revisionism, where it is mandated, but it is hard to shake the thought that there is something suspicious about a position that decouples the perception-action interface from the sensorimotor interface. And it is just such a decoupling that Clark and Chalmers seem to be recommending. After all, no challenge is offered to the thought that the bartender uses his visual sense to access the shapes of the glasses, and it is hard to know what sort of argument could convince us that this is not what is going on. So, if these genuinely sensory events take place inside the cognitive system, and if the perception-action interface is located at the border between mind and world, and if ExC is true, then the perception-action interface must have been decoupled from the sensorimotor interface. Under these circumstances, it strikes me that the revision to our intuitive framework that is demanded of us by the defence of ExC under consideration should at least make us feel uncomfortable. What we certainly don’t have, anyway, is a clear victory for ExC. So where do we go from here?

4. Extended Sensing

Let’s say that we hang on to the thought that the perception-action interface and the sensorimotor interface move around in tandem, and that together they locate the border between mind and world. Is there anything which might give us reason to conclude that the mind-world border, so conceived, is sometimes stationed beyond the skin, meaning that ExC is true? The answer, I think, is yes. Consider cases of what is known as *sensory substitution*. This phenomenon occurs when technological augmentation enables one sensory modality to support the kind of environmental access and interaction ordinarily supported by a different sensory modality. The seminal work in this area is Paul Bach y Rita’s research on tactile-vision sensory

substitution (henceforth TVSS).¹¹ In this work, blind subjects were equipped with a head- or shoulder-mounted camera that conveyed information, from video images, via the activation of an array of vibrators located on the subject's back, abdomen or thigh. After a short period of adaptation, those TVSS subjects who actively controlled the information received, either by manipulating their bodies or by manipulating the camera, were able to make reliable judgments about things such as the number, relative size and position of distal objects in three-dimensional space, and perform actions such as reaching out and picking up objects. TVSS subjects have also been successful at making perceptual judgments involving effects such as looming and object occlusion, and (this time with image-sourced information transmitted via vibrators on the tongue) have reported experiencing illusory movement effects such as the waterfall illusion.¹² TVSS is not the only form of sensory substitution. For example, in auditory-vision substitution,¹³ video images from a camera are converted into sounds (e.g. the vertical positions of pixels may be correlated with different audio frequencies) and conveyed to the subject via headphones. Again after short periods of adaptation, subjects equipped with such devices have been able to localize, and to recognize the shapes of, distal objects in three-dimensional space.

There are some challenging empirical and conceptual issues raised by sensory substitution, and, to clear the way for the lessons that I want to draw from the phenomenon, we need to remark on two of them. First, given that TVSS subjects are perceptually sensitive to distal objects, and because of the phenomenal reports supplied by these subjects, one might wonder whether one ought to say that the perceptual access in question is a form of vision or, alternatively, that it is 'simply' vision-like. This threatens to plunge us headlong into a murky debate over how to individuate the senses.¹⁴ For present purposes, however, we can safely ignore this debate, since we need only the weakest substantive claim in the vicinity, namely that sensory substitution devices support a mode of sensory access to the world that is not the same as the mode of the proximal stimulation that they employ. So, in the case of TVSS, the claim would be that even though the proximal stimuli in play are

¹¹ See, canonically, Paul Bach-y-Rita, *Brain Mechanisms in Sensory Substitution* (New York: Academic Press, 1972). For a more recent treatment, see e.g. Paul Bach-y-Rita and Stephen W. Kercel, 'Sensory Substitution and Augmentation: Incorporating Humans-in-the-Loop', *Intellectica*, 2: 35 (2002), 287-297.

¹² If one stares at a waterfall for a while, and then looks at any stationary rocks at the side of the waterfall, the rocks will appear to be moving upwards.

¹³ Malika Auvray, Sylvain Hannequin, Charles Lenay, and Kevin O'Regan, 'There is Something out There: Distal Attribution in Sensory Substitution, Twenty Years Later', *Journal of Integrative Neuroscience* 4 (2005), 505-521.

¹⁴ For a fresh take on this difficult issue, see Fiona Macpherson, 'Taxonomising the Senses', *Philosophical Studies* 153: 1 (2011), 123-142.

tactile in character, the observed perceptual sensitivity is not correctly categorized as one of touch, even if it isn't vision. Given the distal character of the observed perceptual access, and given that touch is a way of accessing external elements by coming into physical contact with them, it certainly seems uncontroversial to say this much. And we need this much because anything less than a technology-driven change from one perceptual modality to another is unlikely to send us reaching for ExC. Consider, for example, the fact that poor eyesight may be improved by putting on the right spectacles. This is a case of a technology-driven transformation in perceptual performance *within* a single sensory modality. But spectacles are most naturally interpreted as modifying the causal inputs to the relevant psychological machinery (by bending light rays), rather than as part of that machinery. After all, the proximal sensory channel here is already set up so as to support distal perceptual sensitivity prior to the augmenting technology being added. By contrast, in, say, TVSS, the proximal sensory channel is not already set up so as to support distal perceptual sensitivity. Such sensitivity becomes possible via that channel only once the augmenting technology is added. The contribution made by the technology is thus more likely to motivate the claim that the psychological capacity itself is partly realized by the augmenting technology.

The second issue on which I need to comment is that thinkers such as Alva Noë, Mirko Farina and Julian Kiverstein have argued that sensory substitution provides the raw material for an argument in support of what we can call the hypothesis of extended phenomenal consciousness.¹⁵ For the uninitiated, phenomenal consciousness is the *what-it's-like-ness* of experience, such as what it's like for me to see green, taste spaetzle, or, as I'm doing right now, listen to Public Service Broadcasting. If the hypothesis of extended phenomenal consciousness is true, then although the machinery that realizes conscious phenomenal experience includes, and maybe necessarily includes, neural elements, it is not restricted to such elements. There will be cases in which that machinery additionally includes not only non-neural bodily elements, but also elements located beyond the skull and skin. Whether extended phenomenal consciousness ever occurs, and even if it does whether sensory substitution provides supporting evidence for it, are vexed questions.¹⁶ However, this is another debate that, for present purposes, we can

¹⁵ See: Alva Noë, *Action in Perception* (Cambridge, Mass.: MIT Press, 2004); Alva Noë, *Out of Our Heads: Why you are not your Brain, and other lessons from the Biology of Consciousness* (New York: Hill and Wang, 2009); Julian Kiverstein and Mirko Farina, 'Do Sensory Substitution Devices Extend the Conscious Mind?', in Fabio Paglieri ed. *Consciousness in Interaction: the Role of the Natural and Social Context in Shaping Consciousness* (Amsterdam: John Benjamins, 2012), pp.19-40.

¹⁶ For critical discussion, see e.g.: Andy Clark, 'Spreading the Joy? Why the Machinery of Consciousness is (probably) still in the Head', *Mind* 118: 472 (2009), 963-993; Michael Wheeler, 'Not What it's Like but Where it's Like: Phenomenal

happily put aside, because as long as there is a distinction *within* the class of the psychological, between, on the one hand, conscious mental states and processes, and, on the other, unconscious or nonconscious mental states and processes, it is, in principle, possible to develop ExC as a claim about the latter while simultaneously holding that phenomenal consciousness is realized entirely in the brain. If this is the right way to divide things up, then while ExC is true (as a restricted claim about unconscious or nonconscious mental states and processes), the hypothesis of extended phenomenal consciousness is false. In the present context, what this means is that even if the psychological machinery that realizes the conscious perceptual experiences enjoyed by TVSS subjects is wholly neural, the underlying psychological machinery that realizes the process by which one sensory modality supports the kind of environmental access and interaction ordinarily supported by a different sensory modality may include all sorts of processing routines of which the subject is not conscious, and thus may be extended so as to encompass the skin-external augmenting technology. There's nothing especially mysterious about this disassociation. After all, my conscious registering of some item of information recalled from my organic memory may occur in some particular group of neurons, but that conscious registering may be supported by access and storage mechanisms elsewhere in the brain, mechanisms of which I remain unaware, even though they partly realize the psychological process of remembering that's going on.

We are now in a position to make what is, with respect to the received view in and around cognitive science, a bold but exciting claim, one that I cannot hope to defend fully here. Given that the shift in perceptual modality that sensory substitution involves (e.g. from touch to a modality that is certainly not touch and which bears at least some similarities to vision) depends crucially on the contribution of the skin-external augmenting technology, and given that we have resolved not to be distracted by questions regarding the locations of the vehicles of the subject's conscious experience, it seems that sensory substitution is at least in the right ballpark to count as an example of extended sensing, and thus, if we maintain the co-location of sensing and perception, of extended perception. If this is the correct interpretation of sensory substitution, then the advocate of ExC can agree with Chalmers that the mind-world border ought to be stationed at the perception-action interface, because, according to the proposed model, it is that very interface that has been shifted outwards, beyond the skin, to the periphery of the augmenting technology. And if the border between mind and world has been relocated to a point beyond the skin, then ExC is true.

Consciousness, Sensory Substitution and the Extended Mind', *Journal of Consciousness Studies* 22: 3-4 (2015), 129-47; Michael Wheeler, 'Extended Consciousness: an Interim Report', *The Southern Journal of Philosophy* 53, Spindel Supplement (2015), 155-75.

5. Beyond the Container

I have just suggested that Chalmers' objection to ExC may be met, because sensory substitution should be understood as an example of extended sensing. But this is not the final resting place of my argument. To bring that final destination into view, we need to reveal an unexpectedly conservative facet of ExC, as presented so far. To achieve this, let's start not with ExC, but with our old friend vehicle internalism. Even though vehicle internalism has it that cognition takes place in the brain, which makes it metaphysically opposed to substance dualism, and even though Descartes' philosophy is perhaps the historical pinnacle of substance dualism, nevertheless the fact is that vehicle internalism, as standardly developed, has a decidedly Cartesian heritage. As Jennifer Hornsby points out:

...philosophers of mind have come to see Cartesian dualism as the great enemy, but have underestimated what they have to contend with. Taking the putatively immaterial character of minds to create the only problem that there is for Descartes' account, they marry up the picture of the person with the picture of her brain and settle for a view of mind which, though material in its (cranial) substance, is Cartesian in its essence.¹⁷

There are a number of dimensions along which the Cartesian legacy of orthodox vehicle internalism may be laid bare,¹⁸ but here I want to focus on an aspect that intersects with our present interest in borders. In the movement of thought through which philosophers, psychologists, neuroscientists and others took the Cartesian mind and shoehorned it into the body/brain, the seductive idea that the perception-action interface is the border between mind and world became allied with a view of the edge of the body/brain as a kind of physical container within which the mechanisms of mind are housed. Of course, the container is not impermeable: as we have seen, structure flows into the container via perception (so that inner thought comes to reflect the external world) and out of the container via action (so that the external world comes to reflect our inner desires). Nevertheless, it encloses, in a fixed way, the material mechanisms of thought, just as the realm of immaterial substance 'encloses', in a fixed way, the immaterial mechanisms of thought in Cartesian substance dualism.

¹⁷ Jennifer Hornsby, 'Physicalist Thinking and Conceptions of Behaviour', in Philip Pettit and John McDowell eds., *Subject, Thought, and Context* (Oxford: Oxford University Press, 1986), pp.114-15

¹⁸ For example, in the principle that perception, thought and action are conceptually and functionally distinct, and in the claim that intelligent action is always governed by reason. See Michael Wheeler, *Reconstructing the Cognitive World: the Next Step* (Cambridge, Mass.: MIT Press, 2005).

It is at this juncture that a profoundly conservative dimension of ExC, as we have encountered it so far, is revealed. Consider: from what we have just seen, how much damage is caused to the Cartesian cognitive container by the arrival of the extended mind? Plausibly, the answer is ‘not much’. The container has become enlarged, of course, so that the border between mind and world is no longer stationed at the edge of the brain or body, but at the other side of a notebook, smartphone, arrangement of glasses, or TVSS device. Nevertheless (this line of reasoning goes), that border still acts as container for our psychological mechanisms in a remarkably familiar way. Of course, it’s not entirely familiar. The cognitive container is now one that gets bigger and smaller. For example, expert users of TVSS devices routinely experience distal objects and states of affairs in the world via the substituting (not-touch) perceptual modality. However, with conscious, deliberate effort, those users are sometimes able to refocus their conscious attention, so that they experience the *proximal* sensory inputs, that is, tactile experiences of the vibrations taking place on the surface of their skin.¹⁹ Here one should not be distracted by the fact that we have resolved not to think of sensory substitution as supporting the hypothesis of extended phenomenal consciousness. The shifting of conscious attention here ‘merely’ provides evidence that the mind-world border, as located by the functionally defined sensorimotor (perception-action) boundary, may shift dynamically, between a position located at the skin and one located beyond it. The resulting picture is thus one of a dynamic container, one that grows and shrinks as skin-external elements become incorporated into, and divested from, the cognitive system, during its ongoing activity. Still, a rubbery container remains a container, something that confines, encloses, and constrains. And, as framings of the mental go, this simply doesn’t do justice to the fact that out our cognition-realizing materiality doesn’t merely provide a channel by which we achieve epistemic and practical access to a world beyond us, it also plays a conceptually prior role in constructively founding both (a) the world, as an intelligible place that may be accessed epistemically and practically by us, and (b) us, as beings who are essentially of that world. In short, we need to get beyond the container model, if we are to turn sensing into sense-making.

Heidegger has his eye on this very prize when he rejects (what is tantamount to) the container model, with the following words: ‘the perceiving of what is known is not a process of returning with one’s booty to the ‘cabinet’ of consciousness after one has gone out and grasped it’.²⁰ And indeed it is Heidegger who, I think, offers us some of the conceptual resources that might enable cognition to escape from the Cartesian container, in that he supplies us with an account of spatial borders in general that may be used to generate a new account of the mind-world border in particular.

¹⁹ See Noë, *Out of our Heads*, p.62.

²⁰ Martin Heidegger, *Being and Time*, translated by John Macquarrie and Edward Robinson (Oxford: Blackwell, 1927/1962), p.89.

According to Heidegger, spatial borders are not limits or enclosures. Rather, they are productive events – what he calls *horizons* – that establish sites for sense-making. As he puts it:

A space is something that has been made room for, something that has been freed, namely within a boundary... A boundary is not that at which something stops but, as the Greeks recognized, the boundary is that from which something *begins its essential unfolding*. That is why the concept is that of *horismos*, that is, the horizon, the boundary.²¹

Straightaway, any Heidegger scholars out there will no doubt be clamouring to complain that when Heidegger conceptualizes space in this way, he is concerned not with mathematically describable physical space – the kind of space with which we have been concerned up to now – but rather with the space that is relevant to sense-making, or what we might call *existential space*, the collection of spaces in which we dwell. Heidegger calls such existential spaces *locales*. A bridge, for example, brings a locale into existence by establishing a ‘passage that crosses’.²² Heidegger claims explicitly that physical space is an abstraction from existential space, and furthermore that once one thinks in terms of physical space, one has lost all sight of locales.²³ But it seems to me that this final move ought to be challenged, because it’s hard to resist the thought that locales are realized in physical space. Indeed, Heidegger himself describes the locale founded by the bridge as encompassing the banks of the stream (while also transforming them into banks that *lies across from each other*) as well as the surrounding earth (while transforming it into *landscape*). The best way to express this point, I think, is to say that the borders of locales are both mathematically and existentially spatial.²⁴

²¹ Heidegger, ‘Building. Dwelling. Thinking’, p.356.

²² Heidegger, ‘Building. Dwelling. Thinking’, p.355. In truth, my brief gloss on the notion of a locale barely hints at the complexity of Heidegger’s account, which involves what he calls the ‘gathering of the fourfold’ (earth, sky, divinities and mortals). I do not have the space here to explore this intricate notion. For my interpretation, see Michael Wheeler, ‘Martin Heidegger’, *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta ed., section 3.

<http://plato.stanford.edu/entries/heidegger/>, last accessed 04/09/2015.

²³ Heidegger, ‘Building. Dwelling. Thinking’, p.357.

²⁴ For critical commentaries on Heidegger’s account of space that go well beyond the very swift criticism that I offer here, see e.g.: Hubert L. Dreyfus, *Being-in-the-World: a Commentary on Heidegger’s Being and Time, Division 1* (Cambridge, Mass.: MIT Press, 1991), chapter 7; Jeff Malpas, ‘Heidegger, Space and World’, in Julian Kiverstein and Michael Wheeler, eds., *Heidegger and Cognitive Science* (Basingstoke: Palgrave Macmillan, 2012). These commentaries concentrate mostly on the account of space

Of course, minds – even spatial ones – are not locales. Indeed, because the contours of how we find ourselves in the world will partly be determined by the psychological structures that we bring to the sense-making party, it seems that minds are implicated, at a fundamental level, in the very founding of locales. But this fact does not make a consideration of locales irrelevant to our present concerns. We are pursuing the idea that a Heideggerian account of spatial borders might help us to free ExC from the clutches of containerism about the mind. Locales give us a clue as to how this might work, since they embody an account of such borders not as enclosing or limiting surrounds, but as horizons – productive events in histories of sense-making. To see what matters, we need to take two steps. First, it is a feature of Heidegger’s account of human sense-making that mind and world co-emerge as interdependent functions of each other.²⁵ So, just as minds will be implicated, at a fundamental level, in the very founding of locales, so locales will be implicated, at a fundamental level, in the very founding of minds. Minds, too, are productive events in historical flows of sense-making. Second, if we keep this co-emergence picture in view, Heidegger’s work is arguably in the nearby background when Christine Battersby recommends a shift from containers to event-horizons in our conceptualization of the edge of the body. Inspired also by the science of self-organizing dynamical systems, in which global form arises not from external or executive control, but in an autonomous, emergent fashion from nonlinear causal interactions between lower-level elements, Battersby writes as follows:

[F]orms are not fixed things, but temporary attestations in continuous metastable flows, potentialities or evolutionary events. If we think about boundaries, then... the boundaries of bodies need not be thought as the edges of ‘three-dimensional containers into which we put certain things... and out of which other things emerged’. The boundary of the body can also be thought as an event-horizon, in which one form (myself) meets its potentiality for transforming itself into another form or forms (the not-self). Such a body-boundary entails neither containment of internal forces nor repulsion of/protection against external forces. Those who are aware of themselves as centred ‘inside’ an insulated container – free from contamination by the threatening other which is located on the ‘outside’ – are captured by an illusion generated by the mechanisms of ego-

given by Heidegger in *Being and Time*, but the criticisms lodged plausibly generalize to the account he gives in ‘Building, Dwelling, Thinking’.

²⁵ Put *very* crudely, the picture in *Being and Time* is that Heideggerian worlds are structured by the interests and concerns of socially embedded human beings, while the activities of individual human beings are to be interpreted in terms of the structures and possibilities that such beings inherit within such worlds.

protection, as well as by spatial models inherited from a classical science which is now no longer compelling.²⁶

For Battersby, then, the bodily boundary is not an inert territorial wall between self and other that is traversed or breached, but rather an active site of production – that from which the self, as a relational form, begins an unfolding into (a realization of) possibilities. To say that the self is contained within the bodily boundary is to ignore the fact that the self is essentially a function of that boundary, a boundary which is itself only a more or less temporary, more or less stable structure in a dynamic and historically evolving flow.

In a manner, all I wish to add to Battersby's (if I am right) Heidegger-influenced horizontal account of the edge of the organic body is the claim that a similar horizontal approach may be profitably adopted toward yet another spatial boundary, namely the edge of thought. That too is a productive event that allows the mind to unfold in many different forms, some extended and some not. So neither skull nor skin constitutes a container for the mind. In fact, nothing does. Moreover, there is no single, fixed edge of thought coinciding with the skull or skin. Thought has many edges, or, if you prefer, it has a moving, shifting edge. Put another way, there is no single, fixed, mind-world border coinciding with the organic sensorimotor interface. In different circumstances, that border is founded, refounded and refounded again, as skin-external elements become incorporated into, and divested from, the cognitive system, during its ongoing activity. This amounts to what John Sutton has dubbed a 'deterritorialized' view of cognition, a view 'which dissolves individuals [and individual minds] into peculiar loci of coordination and coalescence among multiple structured media' and which analyzes the 'boundaries [between inner and outer, natural and artificial] as hard-won and fragile developmental and cultural achievements, always open to renegotiation'²⁷. Importantly, this deterritorialized picture does not require us to struggle against the intuitive pull of the idea that the mind-world border is located at the perception-action interface, for, as we have seen, that very interface, reconceived as a productive event-horizon on the model of spatial borders recommended by Heidegger and Battersby, is itself a dynamic, moving structure. So, where is the love? Sometimes it's over here, sometimes it's over there, and sometimes it really is two metres to the left of the door, just above the sofa, with a volume about the size of a football.²⁸

²⁶ Christine Battersby, *The Phenomenal Woman: Feminist Metaphysics and the Patterns of Identity* (London: Routledge, 1998), p.52.

²⁷ Sutton, 'Exograms and Interdisciplinarity', p.213.

²⁸ Thanks to two anonymous referees whose comments on an earlier version of this paper enabled me to make improvements.