EasyChair Preprint

№ 15366

Visceral Extension and the Beast-Machine: Toward a Neuropsychological Criterion for the Extended Self

Joseph Neisser

November 5, 2024

Visceral extension and the beast-machine: Toward a neuropsychological criterion for the extended self

Joseph Neisser[a]

[a] *Department of Philosophy, Grinnell College, Grinnell, IA, USA*

ORCID ID: 0000-0002-2909-5029

Abstract. When does a cognitive system distributed across brain, body, and world constitute an extended *self* or person? If cognition extends "beyond skin and skull," then does the self also extend? Or is cognitive labor simply outsourced to an external device that is not part of the self? More generally, what does reflection on technologies of cognitive enhancement suggest about the nature of the self and its relation to the brain and body? Drawing on Anil Seth's concept of the self as a "beast-machine," I argue that although (some form of) the Parity Principle may plausibly provide a criterion for cognitive extension, this is not sufficient to establish an extended self, and that the self is not co-extensive with cognition.

Keywords: Extended Cognition, Embodiment, Self, Predictive Processing, Consciousness

When does a cognitive system distributed across brain, body, and world constitute an extended *self* or person? If cognition extends "beyond skin and skull,"[1] then does the self also extend? Or is cognitive labor simply outsourced to an external device that is not part of the self? More generally, what does reflection on technologies of cognitive enhancement suggest about the nature of the self and its relation to the brain and body? Drawing on Anil Seth's [1] concept of the self as a "beast-machine," I argue that while the received form of the Parity Principle plausibly provides a criterion for cognitive extension, it is not sufficient to establish an extended self. That is, the self is not co-extensive with cognition. Nor is it co-extensive with consciousness. The embodied self runs deeper than consciousness; the lived body is not simply the body as it appears in consciousness. I suggest a conception of extended selves rooted in neural mechanisms for proprioception and interoception.

First, I'll quickly rehearse the background for extended cognition which, doubtless, is familiar to many readers. Begin with the *Parity Principle*:

> If, as we confront some task, a part of the world functions as a process which, were it to go on 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. ([2], p.8)

Next, add the so-called *Glue & Trust Criteria*:[2]

> 1. That the external resource be reliably available and typically invoked.
> 2. That any information thus retrieved be more-or-less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example). It should be deemed about as trustworthy as something retrieved clearly from biological memory.
> 3. That information contained in the resource should be easily accessible as and when required. ([5], pp.79)

---

1 The phrase "beyond skin and skull" is due to [2].

2 The phrase "glue & trust" is from [3]. See also [4].

By appeal to the Parity Principle, duly supplemented by the glue & trust criteria, it is possible to construct arguments showing that notebooks, video game controllers, smart phones, and gps devices can become parts of an integrated cognitive system that extends beyond skin and skull. These extended systems perform labor for which we would give cognitive credit, had a person done it in their heads. In that sense, thought itself is extended and enhanced.

In a later discussion of the metaphysical implications of the Parity Principle, David Chalmers [6] distinguishes between two kinds of extended cognition, *circuit extension* and *sensorimotor extension* (see also [7]). I next introduce these and then suggest a third possibility which crosscuts Chalmers' distinction: *visceral extension.*

Circuit extension is exemplified by the case of Diva, a human who suffers minor brain damage and loses some specific cognitive ability, say, arithmetic [8]. Diva has an external silicon circuit connected wirelessly to her brain, restoring the original functions without loss. By parity reasoning, then, Diva's cognition is now extended in a new way, supervening on the chip-brain circuitry rather than just the brain. Even the critics of extended cognition find this conclusion undeniable (e.g., [9]). But *sensorimotor extension* is more interesting and more difficult to accept. In sensorimotor extension, the glue & trust criteria are met via person-level perception and action rather than by subpersonal circuitry. Amnestic Otto accesses information in his trusty notebook by *reading* it, and the savvy Tetris player finds the correct orientation of an image by *turning* it on the screen with a dependable clicker [2]. Parity reasoning (again supplemented by Glue & Trust) shows that reading the notebook is a case of remembering, and turning the images is a case of mental rotation.

Thus, sensorimotor extension demonstrates that one's cognitive processes can be outside of *oneself* in surprising ways. Chalmers eventually settles on a revision of the extended mind thesis: A subject's cognitive processes and mental states can be partly constituted by entities that are external to the subject, in virtue of the subject's sensorimotor interaction with these entities ([6], p.7). The upshot is that cognitive processes can be external, not just to the skin and skull, but *to the subject whose cognitions they are*. Chalmers explains that the reason that sensorimotor extension is controversial, while circuit extension is not, is that it violates a fundamental assumption of cognitive psychology, namely, that cognition takes place *between* perception and action. If the thesis of sensorimotor extension is correct, then a core idea of cognitive psychology may turn out to be false. In contrast, the extended circuit thesis is just another piece of banal functionalism. As Chalmers puts it, the thesis of sensorimotor extension "… does not just overthrow the hegemony of skin and skull as boundaries for cognition ... It also overthrows the hegemony of perception and action." ([6], p. 9)

As it stands, Chalmers' distinction between circuit extension and sensorimotor extension suggests a simple application to the question about extended selves. Circuit extensions might be self-extensions because they are functionally "between" the inputs and outputs –perception and action – that constitute the bounds of an experiencing subject. Sensorimotor extensions, meanwhile, would *not* be part of the extended self since they would not be between perception and action. Because one must interact with them on a personal level, they are not part of one's person. But this is too quick. Chalmers speaks of "the subject," not the self or person. Often, writers do speak of "the subject" in a way which suggests this equivalence with the self [10]. But subjectivity corresponds to the first-person *perspective*, which can be enacted in conscious creatures that lack a self [11]. Subjectivity in this sense is partly constitutive of consciousness,

and it is to the possibility of *extended consciousness* that Chalmers next turns his attention, arguing that the parity cases are not cases of extended consciousness. Roughly, he holds that perception and action still constitute the bounds of consciousness, even when cognition transcends these bounds via sensorimotor extension. In this way, Chalmers' distinction between sensorimotor extension and circuit extension recapitulates the distinction between consciousness and cognition.

But the self crosscuts the distinction between consciousness and cognition. Pace any doctrinaire Husserlian or Sartrean phenomenology, the self cannot be straightforwardly identified with conscious subjectivity. If there are non-conscious embodied or neural dimensions of the self, then the bounds that constrain consciousness may not constrain the self. Like cognition, the self may extend beyond and beneath consciousness. But unlike cognition, the self is no mere matter of representation or information processing, because information processing by itself is also insufficient to constitute a self. Whatever an extended self is, it is not just an extension of one's thought but of one's *embodiment*. A self is not just a "thinking thing" (though it is that). It is also a particular kind of being in the world, a way of being that distinguishes itself from its world and carves itself out from the world. So, an extended self is an extended, thinking body that is somehow *lived*. Just as the self is not merely identical with whatever one thinks about or values, the self is also not just whatever physical substrate realizes one's cognition. To extend the self is to become entwined or entangled with an assemblage in a distinctive way.

If the self is embodied and lived in this way, then the Parity Principle cannot be sufficient, even when supplemented by Glue & Trust, to ground the attribution of an extended self (though it may be necessary). A new category is required, which I call *visceral extension.* To evoke this idea, I turn to Anil Seth's image of the self as a *beast-machine* [1].[3] The phrase "beast-machine" is meant, in part, to evoke our basic sense of *being an embodied organism* ([1], pp.250). Roughly, Seth's thought is that embodied selfhood may be grounded in a particular kind of predictive processing, distinct from that which underlies perception of the external world.

Proponents of predictive processing hold that the brain is fundamentally a Bayesian machine. Paradigmatically applied to explain perception, the theory holds that perception consists in the brain's expectations about the next sensory inputs - an expectation which functions as an inference about the external causes of stimuli. Seth [1] extends the framework to speculate that feelings of selfhood are the result of a subtly different, non-perceptual kind of predictive processing, which he calls *control-oriented active inference*. While perception of the external world involves *epistemic* inference about the causes of stimulation, feelings of the internal milieu may arise from control-oriented inferences which regulate homeostasis ([1], pp. 242–243).

Seth's approach could be helpful for formulating an appropriately embodied and genuinely neuropsychological criterion for the extended self. Such a criterion would ask whether the extending technology – suitably glued & trusted – is coupled with the neural mechanisms that maintain the embodied self – i.e., the beast-machine. The brain may interface with technology in multiple ways, including either via epistemic or control-oriented mechanisms. In the latter case, it may be plausible to suppose that the extension has become part of the self. In this case, the relation to the device would be properly *interoceptive*. It would then be a further

[3] Seth draws inspiration for the beast machine from materialist philosopher La Mettrie [12].

question whether this is the only possible basis for an extended self. If so, the criterion may be too restrictive – the technology might then have to be governed by specific and dedicated neural circuits that are not particularly plastic in their range of functions (e.g., brainstem regulation of heart rate), and this kind of brain-technology interface may be limited to certain very specific medical contexts (compare the case of Diva). But if the control-oriented predictive processing for the extended self can include *proprioceptive* mechanisms – mechanisms that perceive, predict, and control the position of the body – then more technological extensions might genuinely become part of the lived body, part of the beast-machine. In effect, proprioception could be a sub-type of sensorimotor extension in which the embodied self *does* overflow the boundaries perception and action set by Chalmers. This idea is not without problems. Seth's framework is designed around interoception. Proprioception would be excluded on grounds that it is perceptual rather than homeostatic. So, if visceral extension is our sole criterion, then technologically extended selves will be rare indeed. This may in fact be he best way forward – it may be best to deny that popular cases of extended cognition are also cases of the extended self. I am sympathetic to this thought. But it may also be too restrictive. It remains highly intuitive that proprioception is, properly speaking, *self*-perception. If so, and if proprioception is not illusory, then there are grounds for including the neural mechanisms of proprioception in the beast-machinery. Certainly, this does not settle the matter. But it points toward a slightly more permissive version of Seth's framework.

I conclude that technology becomes a proper part of an embodied self when it is anchored in the control-oriented predictive processing mechanisms in the mid-line neural systems implicated in self-regulation, self-control, and self-perception. Thus, the way to *visceral extension* of the embodied self leads through the default mode network.

If the above is on the right track, there are a range of implications for emergent questions about personal integrity, autonomy, and privacy in the age of enhanced reality. For cases of extended cognition *without* an extended self, it may be that existing policy and case law might be applied quite directly. Extended cognition in the absence of self-extension might be comparable to other kinds of personal data – perhaps at one end of a spectrum of such data. Genuinely extended selves, however, would have a different status, and infringements by others might be best understood as violent assault.

Interesting speculative questions also ensue. The nature and even the existence of the self is not well understood. If we can identify the conditions under which extended selves really do and do not arise, how does this affect the wider debates about the self? How might we better conceptualize a "post-Cartesian" kind of being in the world, and contribute to the widely announced project of rethinking the humanistic tradition? When the prototypical human subject is just one possible formation of the self, the effect is to "decenter the human" and to suggest a continuum between human, animal, and technological forms of cognition ([12], p.182). An account of the extended self can play a role in transforming, while retaining contact with, traditional concepts like agency, responsibility, and integrity. In the world of technological entanglement these categories will not be simply jettisoned, They will be retrofitted.

Notoriously, the late Daniel Dennett asked, “Where am I?” [14]. He devised a series of cases involving a telepresence robot of a kind that was strictly science fiction at the time.[4] His characteristically rhetorical discussion pointed to a deflationary, possibly eliminativist conception in which the self is either identified with the brain or turns out to be a narrative illusion (see also [15]). But perhaps there really are embodied selves that can extend across the brain, animal body, and technology?

References

[1] Seth, Anil K. (2019). Being a beast machine: The origins of selfhood in control- oriented interoceptive inference. In Columbo, Irvine & Stapleton (eds.), *Andy Clark and His Critics*. Oxford University Press: 238–253. https://doi.org/10.1093/oso/9780190662813.001.0001

[2] Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1*)*: 7–19.

[3] Clark, A. (2010). Coupling, constitution, and the cognitive kind: A reply to Adams and Aizaway. In R. Menary (ed.), The Extended Mind. Cambridge, MA: MIT Press, pp. 81–100.

[4] Gallagher, S. (2023). *Embodied and Enactive approaches to Cognition.* Cambridge University Press.

[5] Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Ex-tension*. OUP USA.

[6] Chalmers, D, (2019). Extended cognition and extended consciousness. In Columbo, Irvine, & Stapleton (eds.), *Andy Clark and His Critics*. Oxford University Press: 9–20. https://doi.org/10.1093/oso/9780190662813.001.0001

[7] Parsons, T. & Neisser, J. (2024). Intersecting frontiers of extended reality and neuropsychology. In Kissel, A. & Ramirez, E. (Eds), *Exploring Extended Realities: Metaphysical, Psychological, and Ethical Challenges.* Routledge Press.

[8] Clark, A. (2009). Letter to the editor. *London Review of Books 31.6* (March 26). https://www.lrb.co.uk/the-paper/v31/n06/letters

[9] Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*(1): 43–64.

[10] Neisser, J. (2008). Subjectivity and the limits and narrative. *Journal of Consciousness Studies* 15, no.2, 51-66.

[11] Neisser, J. (2015). *The Science of Subjectivity*. Palgrave Macmillan Press.

[12] La Mettri, J.O. (1748). *L’homme Machine*. Leiden : Luzac

[13] Hayles, N. K. (2017). *Unthought: The Power of the Cognitive Nonconscious*. Chicago University Press.

[14] Dennett, D.C. (1978). Where Am I? In *Brainstorms: Philosophical Essays on Mind and Psychology.* Montgomery, VT: Bradford Books: 310–323. https://mitpress.mit.edu/9780262540377/brainstorms/

[15] Dennett, D.C. (1992). The self as a center of narrative gravity. In F. Kessel, P. Cole and D. Johnson, eds, *Self and Consciousness: Multiple Perspectives*, Hillsdale, NJ: Erlbaum.

[4] See [7] for discussion of the extended self and a contemporary version of the telepresence robot – one that is no longer science fiction.