

Extended animal cognition

Marco Facchin, Ph.D.

[corresponding author]

Research Fellow

Universiteit Antwerpen

Centre for Philosophical Psychology

Prinstraat 13, Antwerp 2000

marco.facchin@uantwerpen.be

marco.facchin.marco.facchin@gmail.com

<https://orcid.org/0000-0001-5753-987>

Giulia Leonetti, MA.

Ph.D. Candidate

Istituto Universitario di Studi Superiori IUSS

Pavia

Linguistics & Philosophy IUSS Center

(L&PIC), Department of Human and Life

Sciences; Pavia, Italy

Palazzo del Broletto, Piazza della Vittoria n.

15, 27100, Pavia

giulia.leonetti@iusspavia.it

<https://orcid.org/0000-0002-2422-6755>

Abstract:

According to the extended cognition thesis, an agent's cognitive system can sometimes include extracerebral components amongst its physical constituents. Here, we show that such a view of cognition has an unjustifiably anthropocentric focus, for it depicts cognitive extensions as a human-only affair. In contrast, we will argue that if human cognition extends, then the cognition of many non-human animals extends too, for many non-human animals rely on the same cognition-extending strategies humans rely on. To substantiate this claim, we will proceed as follows. First (§1), we will introduce the extended cognition thesis, exposing its anthropocentric bias. Then, we will show that humans and many non-human animals rely on the same cognition-extending strategies. To do so, we will discuss a variety of case studies, including "intrabodily" cognitive extensions such as the spinal cord (§2), the widespread reliance on epistemic actions to solve cognitive tasks (§3) and cases of animal cognitive offloading (§4). We'll then allay some worries our claim might raise (§5) to then conclude the paper (§6).

Keywords: Extended cognition, Epistemic actions, Non-human cognition, Animal Cognition, Cognitive offloading

Declarations: NA

Founding: MF is currently funded by the FWO grant "Towards a globally non-representational theory of the mind" (grant number 1202824N). GL is funded by a Ph.D. grant in Sustainable Development and Climate Change (SU4). This paper and related research have been conducted during and with the support of the Italian national inter-university PhD course in Sustainable Development and Climate change (link: www.phd-sdc.it).

Conflict of Interests/competing interests: The author declares no conflict of interests.

Availability of data and Material: Not applicable.

Code availability: Not applicable.

Authors' contribution: GL ideated the paper. MF structured the arguments and wrote the first draft. Both authors are equally responsible for the final version of the manuscript. Both authors agreed on the authorship order.

Acknowledgements: The authors wish to thank Marco Viola and Sanja Sreckovic for having read and commented upon various early drafts of the manuscript. We also wish to thank the participants of the Rudolf Carnap Lectures in Bochum and the SINE conference in Pisa and Lucca for their stimulating questions. A special thank you to Wouter van Hooydonk for his suggestions concerning human exceptionalism. Thanks also to the two anonymous referees for their stimulating comments

Extended Animal Cognition

Abstract:

According to the extended cognition thesis, an agent's cognitive system can sometimes include extracerebral components amongst its physical constituents. Here, we show that such a view of cognition has an unjustifiably anthropocentric focus, for it depicts cognitive extensions as a human-only affair. In contrast, we will argue that if human cognition extends, then the cognition of many non-human animals extends too, for many non-human animals rely on the same cognition-extending strategies humans rely on. To substantiate this claim, we will proceed as follows. First (§1), we will introduce the extended cognition thesis, exposing its anthropocentric bias. Then, we will show that humans and many non-human animals rely on the same cognition-extending strategies. To do so, we will discuss a variety of case studies, including "intrabodily" cognitive extensions such as the spinal cord (§2), the widespread reliance on epistemic actions to solve cognitive tasks (§3) and cases of animal cognitive offloading (§4). We'll then allay some worries our claim might raise (§5) to then conclude the paper (§6).

Keywords: Extended cognition, Epistemic actions, Non-human cognition, Animal Cognition, Cognitive offloading

1 - Introduction

According to the extended cognition thesis (henceforth EC) an agent's cognitive system is at times partially constituted by elements lying outside the agent's brain (Clark & Chalmers 1998; Clark 2008). If EC is correct, then, cognitive systems include amongst their bits and pieces certain extracerebral components, such as specific bodily actions or certain environmental props.

EC had a sizable impact on the philosophy (see Menary 2010; Colombo, Irvine & Stapleton 2019) and empirical practice (e.g. Vallé-Tourangeau *et al.* 2016; Bocanegra *et al.* 2019) of cognitive science, sparking a heated debate concerning the correct interpretation of EC and its truth. As important as that debate is, we won't engage with it here. We will simply assume that some variant of EC is correct, to focus on an aspect of the debate that we find problematic.

To see that problematic aspect, consider the fairly *standard* formulation of EC given above. It begins by mentioning "cognitive agents", which is pretty generic. In principle, many different system types may be cognitive agents, including humans, non-human animals (henceforth just "animals", for the sake of brevity), plants, robots, and fungi. As formulated above, EC can be applied to *all* these systems. This, of course, does not entail that EC is *true* of all of them: indeed, EC may be false in some or all of these cases. But even so, EC would be false only because it can be *literally and legitimately* applied to each of these systems (eg. Figdor 2018).

However, the standard formulation of EC given above soon gains *at least* a vertebrate-centric focus, as it explicitly mentions *brains*. It thus seems that the scope of EC has been significantly restricted to a small subset of all these systems; namely the subset of animals with brains (eg. Sims & Kiverstein 2022). And, really, if one looks at how EC is *de facto* discussed, the word "brain" seems to refer exclusively to the *human* brain. For, barring a few papers (e.g. MacIver 2009; Japyassú & Laland 2017; Sims & Kiverstein 2022; Parise *et al.* 2020, 2023), EC has only been applied to humans. Just consider the standard, go-to cases invoked when discussing EC: they're button presses when playing video games (Kirsh & Maglio 1994; Clark and Chalmers 1998), smartphones (Chalmers 2008), wearable "intelligent" technologies (Clark 2003; Bach-y-Rita & Kercel 2003; Wheeler 2019a) and, more recently, the internet (Smart 2021). Even "low-tech" cognitive extensions seem exclusively human, as they are typically exemplified by language (Clark 1997) and notational systems (Dutih-Novaes 2012; De Cruz & De Smedt 2013; Malafouris 2013; Menary 2015).

This human-centric focus should not be overstated. As indicated above, some papers have already applied EC to non-human systems. And we find it unlikely that the thinkers mentioned above are really committed to the claim that EC can be applied only to humans. Indeed, it is hard to see how they *could* hold that EC applies only to humans. EC is typically defended on either functionalist (e.g. Clark & Chalmers 1998; Clark 2008) or dynamicist (Chemero 2009; Palermos 2014) grounds, and both these grounds allow (or at least, do nothing to prevent) EC to be applied to animals. Thus, we find it more likely that EC's human-centric focus is due more to the contingent fact that philosophers interested in EC are generally not very interested in animal minds, rather than to a theoretical commitment to keep EC a "human only" claim.

Be as it may, the neglect of animals in the EC debate is problematic. For, the absence of animals does make EC appear as a human-only claim.¹ And once extended cognition appears as a human-only affair, it may become natural and even attractive to think of cognitive extensions as playing a relevant role in our distinctive human nature, and of our human cognitive uniqueness. Here's Andy Clark making precisely this point:

"[...] human beings, I want to convince you, are *natural-born* cyborgs. [...] I do believe that it is the plain and literal truth. I believe, to be clear, that it is above all a SCIENTIFIC truth, a reflection of some deep and important facts about (a whiff of paradox here?) our special, and distinctively HUMAN, nature. Certainly I don't think this tendency toward cognitive hybridization is a modern development. Rather, it is an aspect of our humanity, which is as basic and ancient as the use of speech and which has been extending its territory ever since." (Clark 2003, pp. 3-4; capitalizations in the original; see also Wheeler and Clark 2008; Sterelny 2003 for further examples)

Similar views can also be found in (Sterelny 2003; Wheeler & Clark 2008) and, more cautiously, in (Clarks 2002, p.4).

These exclusively human construals of cognitive extensions, however, are misguided. For, as we shall now claim, cognitive extensions, if they exist, are *not* an exclusively human cognitive trait (c.f. Figdor 2022). More precisely, we will defend the following conditional claim: *if* cognition extends in the human case, then, *ceteris paribus*, it also extends (at least) in various animal cases.

Let us clarify two points concerning our claim. First, in this paper we will only be concerned with extended *cognition*, as opposed to extended *mental states*. Whilst these two types of extension are often conflated in the literature, it is possible - and in some case, necessary - to distinguish between them. When dealing with extended cognition, one is dealing with the extra-cerebral realization of *cognitive processes* - that is, with the extra cerebral realizers of a dynamic, temporally unfolding *process* which may operate on states that are not captured by our folk-psychological lexicon and that may even be sub-personal. In contrast, when dealing with extended mental states, one deals with *states* that are typically well-captured by our folk psychological lexicon and that are personal (see Drayson 2010; Carter *et al.* 2018). So, in this paper, we will deal only with extended *cognitive processes* and the extra-cerebral machinery realizing them. It could be that some of our cases (esp. these discussed in §4) may also support the extension of mental states. But such a claim would be an *additional* claim necessitating a *further* argument, which we can't simply provide here.

Secondly, what do we mean when we talk about *cognitive* processing? Extended cognition famously triggered a heated debate concerning what cognition is (see Adams & Aizawa 2001; Akagi 2018; Varga 2018) - a debate that, to many, has shown we currently lack any non-partisan, widely shared definition or characterization of cognition (see Allen 2017; Favela & Martin 2017; Facchin 2022a; Gough 2023a; Figdor

¹ It may also contribute to a morally problematic invisibilization of animals (on which see van de Brandeler 2004).

2024). We are inclined to agree. This, however, does not mean that we cannot tell apart cognitive processes from non-cognitive ones. We are still able to tell apart *paradigmatic* cognitive processes (processes such as memorizing, extracting information, and internally simulating a scenario) from *paradigmatic* non-cognitive processes (such as vomiting, defecating, having a seizure, and digesting). Similarly, we are still able to tell apart paradigmatically cognitive systems (e.g. humans, animals) from paradigmatically non-cognitive ones (e.g. bricks, clouds). Thus, we can still meaningfully talk about cognition and cognitive systems in reference to these cases. And whilst such a procedure is bound to leave a “gray area” of processes that are neither paradigmatically cognitive nor paradigmatically non-cognitive (is a computer’s “memorizing” data really cognitive? Are computers cognitive systems?), this gray area won’t be a problem here, at least insofar we will deal with processes that are *paradigmatically* cognitive and animals, whose status as cognitive systems is largely accepted.²

Is our defense of animal extended cognition needed? We think it is, for a variety of reasons. First, it counteracts the tendency of the EC literature to neglect animals (and non-human systems) more generally, acting as a powerful reminder of the scope of EC. Secondly, it contributes towards a more accurate assessment of the cognitive differences separating human and animal cognition (eg. Figdor 2021, 2022). Of course, we don’t want to *deny* that there are such differences. We wish, however, to counteract views such as the one Clark (2003) voiced above, according to which such differences are, in important ways, due to the presence or absence of cognitive extensions. Lastly, considering animal cognition as extended encourages researchers to create novel experimental paradigms, aimed at investigating the role of animal cognitive extensions. Such novel experimental paradigms have yielded interesting and unexpected results in the human case (see Vallé-Tourangeau *et al.* 2016; Bocanegra *et al.* 2019), and it seems reasonable to expect similar results in the animal case. Thus, defending the applicability of EC to animals may have positive heuristic results.

Our argumentative strategy is simple. Throughout the paper, we will consider various paradigmatic cases of *human* extended cognition, showing, by relying on relevant empirical data, that they have very close animal analogs: if cognition extends in the human case, then it also extends in the animal case or so, at least, we will argue. Here’s the details of our plan: §2 looks at “intrabodily” cognitive extension, arguing that we’ve good reasons to believe they extend animal cognition too. §3 discusses cognition-extending epistemic actions, highlighting their ubiquity in the animal kingdom. §4 discusses artifact-based cognitive extension in the animal kingdom. §5 deals with some worries our views might raise, and §6 provides a brief conclusion.

2 - Warming up: animal intrabodily cognitive extensions

EC often highlights the cognitive role of external objects - like the role of pen and paper in solving a complex equation. Yet nothing, in the letter of EC, forces cognitive extensions to be outside an agent’s *body*. There are also *intrabodily* cognitive extensions: extra-cerebral constituents of the cognitive machinery which are nevertheless spatially located within the agent’s body. Two such extensions have thus far been identified: the spinal cord (Facchin *et al.* 2021) and the gut-brain axis (Boem *et al.* 2021; 2024). Being less counterintuitive than traditional extra-bodily cognitive extensions, these cases provide optimal warm-up cases to discuss animal EC.

The spinal cord is an anatomically defined neural structure constituting the vertebrate central nervous system. Traditionally, the spinal cord has been thought of as nothing but a “big nerve” allowing for brain-sensory periphery communication and taking care of “hard wired” reflexes (eg. Wolpaw & Tenissen 2001), and this view still is the “standard” view of the spinal cord in neuroscience *at large* (see Kandel *et al.*

² Even philosophers defending a very demanding, narrow and anthropocentric conception of cognition seem to accept that at least some animals are genuine cognitive systems (see, for instance, Adams 2018).

2012 pp. 790-810). This traditional characterization, however, appears increasingly outdated: spinal neuroimaging and cognitive psychology studies revealed that the spinal cord plays a key role in numerous cognitive and affective processes, including: the representation of an agent's "space of action" (e.g. Scandola *et al.* 2016, 2020; Sedda *et al.* 2019), motor imagery (e.g. Fiori *et al.* 2014) and motor learning (Vahdat *et al.* 2015).

The gut-brain axis is a functional axis which includes the autonomic and enteric (i.e. intestinal) nervous systems, the endocrine system, the immune system³ and the central nervous system (spinal cord included). One important, and surprising, component of such an axis is the *gut microbiota*: a community of various types of microorganisms that inhabit an animal's body, entering in a mutually beneficial, symbiotic relation with its host. Importantly, modifications in its composition have been shown to have an important impact in shaping an agent's emotional control and spatial memory (see Foster and Neufeld 2013 for a review).

In both cases, what we observe is a dense, bi-directional and cognitively relevant coupling between the brain with some extra-cerebral component, knitting the two together in a single extended cognitive system (eg. Lamb & Chemero 2018; Palermos 2014), in which activities of the spinal cord and of the gut microbiota are a part of the cognitive processing (see Facchin *et al.* 2021; Boem *et al.* 2021 for an in-depth defense of these claims). But *if* one accepts that the spinal cord or the gut-brain axis are intrabodily cognitive extensions in humans, *then* one has at least two reasons to concede that many animals - perhaps all vertebrates - have similar intrabodily cognitive extensions.

First, human morphology is homologous to the morphology of various vertebrate species, especially other mammalian species. Importantly, homologous structures generally play homologous roles. Sure, there are exceptions: non-human forelimbs can't always grasp, and the simian homolog of Broca's area may not be a language area. But these exceptions are only *partial* exceptions (simian forelimbs can grasp, Broca's simian homolog is still a motor area sensitive to quasi-linguistic regularities, see Rizzolatti & Sinigaglia 2006; Wilson *et al.* 2015) to a *generally valid* rule. Hence, it is extremely unlikely that the spinal cord and the gut microbiota will play a cognitive role *only* in the human case. Indeed, both the gut-brain axis and the spinal cord are evolutionarily ancient, at least if compared with behaviorally modern humans. And it is extremely unlikely that they've acquired *all* their cognitive roles in the single evolutionary step that led to modern humans. Cognitive functions just don't "pop up" this way (eg. Cao 2022). It is thus likely that some such functions have been acquired "along the way", and preserved or re-used through various speciation events. But, then we should expect at least some of these cognitive roles to be played in animals too. Hence, we should expect various animals to have intrabodily cognitive extensions.

Secondly, we have *observed* intrabodily cognitive extensions play the relevant cognitive functions in animal models. Rat models showed that the spinal cord is quite the sophisticated learning system, capable of both operant and instrumental conditioning (Grau 2014). Rat models have also shown the spinal cord being able of some *top-down* control in shock-induced antinociception (see Allen *et al.* 2009 for a nice review). Shock-induced antinociception is a form of stress-induced analgesia, whereby exposure to a mild shock inhibits pain reactivity. Since shock-induced antinociception can be disrupted by distractors, it is thought of it as a top-down attentional (and thus likely cerebral) control effect. And yet, not only shock-induced antinociception is observed in the hindlimbs of spinally transected (paraplegic) rats, but distractors effects are maintained in such conditions too (Meagher *et al.* 1993), demonstrating that top-down attentional control *must* be due to some spinal mechanism. On the fairly uncontested assumption that top-down attentional control is cognitive, the spinal cord of rats has been observed to play a cognitive role. Alterations of the gut microbiota of mice - either letting specific microorganisms colonize the intestine of a mice population or *via* fecal transplant - leads to specific alterations of their exploratory behaviors (Bercik *et al.*

³ For a longer, more recent treatment of the immune system as a cognitive system, see (Gough 2023b).

2011, Neufeld *et al.* 2011a,b) and spatial memory storage (Gareau *et al.* 2011, Wang *et al.* 2015) providing an analogous case for the gut microbiota.

There are thus compelling reasons to believe that *if* the spinal cord and the gut-brain axis are cognitive extensions in humans, *then* they are also cognitive extensions in various animals. What is good for the (metaphorical) goose is good for the (literal) gander.

An objector, however, may contend that cases of intrabodily cognitive extension are not really metaphysically interesting.⁴ One reason why EC metaphysically is interesting is that it puts pressure on the idea that perception and action bound cognition; that is, that cognitive processing starts with perception and ends at action (e.g. Hurley 2001). Metaphysically interesting cases of extended cognition challenge that view, forcing us to see perception-action loops as cogs in the cognitive machinery, in a way that clearly transgresses the boundaries traditionally assigned to minds (Chalmers 2019; Clark 2019). Cases of intrabodily cognitive extension, however, do nothing of that sort, and so do not qualify as *genuine* (i.e. metaphysically relevant) cases of extension. They're only cases of what Chalmers calls "extended circuitry", in which the machinery of cognition is merely spatially extended, leaving unchallenged the traditional, broadly speaking "cartesian" views of the mind.

If Clark and Chalmers are right, then, we've not yet shown any instance of animal extended cognition. The next section addresses this important challenge.

3 - Breaching the sensorimotor boundary: on animal epistemic actions

A way to make sense of Clark (2019) and Chalmers's (2019) claim that in cases of genuine cognitive extension perception-action loops are part of the cognitive processing is by casting such loops as *epistemic actions* (Maglio & Kirsh 1992; Kirsh & Maglio 1994; Kirsh 2019; Clark 2008, 2022). Roughly, epistemic actions are embodied actions that transform informational structures in the environment, so as to facilitate the achievement of certain pragmatic goals. Using a classic example: going to a museum is a "standard" pragmatic action, which brings about a desired goal (e.g. seeing an exhibit). Conversely checking the museum's address on a notebook is an epistemic action, which manipulates environmental sources of information to make the desired goal (i.e. seeing the exhibit) more easily attainable (eg. Clark and Chalmers 1998). Now, if Clark, Chalmers and many others (see Menary 2010, Colombo, Irvine & Stapleton 2019) are right, such actions extend our cognitive processing - they are spinning cogs of our cognitive machinery.

Importantly, epistemic actions *need not* involve the manipulation of artifacts or tools. Clark himself discusses many such "prop-less" epistemic actions, such as bodily movements used to retrieve information "just in time" for its usage (Clark 2008, pp. 11-15), the usage of movement to elicit novel, information-rich, multimodal sensory streams (*idem*, pp. 17-22) and the epistemic role of gesturing (*idem*, pp. 123-126). Here, we present some "prop-less" *animal* epistemic actions. Again, our claim is conditional: if "prop-less" epistemic actions extend human minds, then they extend animal minds too.

As a first example of animal epistemic action, we discuss the case of head-bobbing in birds. Just as humans use saccades to extract information from the environment, so too do some species of birds *via* head-bobbing. In fact, head-bobbing, in addition to being useful for balance stabilization (Theunissen & Troje 2017), also allows for depth perception. Proceeding in order, first we introduce saccades as epistemic actions, and then discuss head-bobbing.

⁴ Yet, as (Facchin *et al.* 2021) and (Boem *et al.* 2021) notice, they might still have important implications concerning the methodology and the conceptual toolbox of the mind sciences and biological sciences.

A well-known fact about our visual fields is that they are not uniform: the area of maximal visual acuity (*fovea*) is surrounded by a peripheral region of diminished visual acuity (eg. Huber-Huber, et al. 2021). Saccadic eye movements are rapid and automatic eye movements that control the position of the fovea, to gather high-quality visual information on various salient objects and features (Liversedge & Findlay 2000; Ibbotson & Kregelberg 2011). These movements often collect such high-quality information just-in-time to control an agent's conduct - for example, allowing us to extract externally stored bits of task-relevant information as the need to consult them arises (O'regan 1992; Clark 2008, pp. 11-15; Rowlands 2010). The same can be said about the depth information obtained by birds thanks to head-bobbing.

Head-bobbing is the rhythmic head movement that characterizes several species of birds, including pigeons (*Columba livia*), ring-billed Gulls, (*Larus delawarensis*) and quails (*Coturnix coturnix*) (eg. Nyakatura & Andrada 2014; Lisney & Troje 2016). Thus, it is a common activity that several species of birds engage in by moving their head back and forward in a rhythmic manner. These head movements are helpful in many of the activities that birds perform every day, from foraging to stabilizing their balance (eg. Fujita 2006, Kral 2003). Crucially, certain studies suggest that head-bobbing is an epistemic action playing a variety of functions within birds' cognitive economy. In particular, the "thrust phase" of head-bobbing (i.e. the phase in which birds push their head forward while keeping it motionless with respect to the vertical axis, see Troje & Frosst 2000) allows birds to extract information about the distance of objects thanks to motion parallax (eg. Lisney & Troje 2016; Kral & 2003). Motion parallax is the phenomenon whereby still objects closer to a moving observer appear faster than still objects further away from the observer. Consider, for example, the different perceived velocities of the objects at various distances when gazed from a moving train, by thrusting their head at controlled velocities. The very same phenomenon allows birds to determine an object's distance (Fux & Eilam 2009). Thanks to motion parallax, it is also possible to learn information about the environment. In addition, birds can also generate motion parallax peering, that is, moving the head on one side and the other while resting. In this way they gather high-quality information about their environment (Pritchard & Healy 2018; Kral 2003). Thus, whereas humans perform saccadic eye movements to gather and extract high-quality information about their environment, birds perform head movements. If the former extends human cognition, the latter should extend bird cognition too.

As a second example, consider how jumping spiders move their whole body in order to see. Jumping spiders are active predators, which stalk and proactively seek their prey. This lifestyle requires sophisticated perceptual capabilities, and in fact, unlike many spiders, jumping spiders have a complex visual system. Their anterior median eyes (i.e. the front-facing pair) have a striking visual resolution, allowing them to form detailed images of their surroundings. All other eyes lack such a rich visual resolution, functioning just as motion detectors. Their visual fields, thus, consist of a central high-resolution area surrounded by a larger low-quality sea of motion detection. As seen above, humans have a similarly structured visual field, and gather environmental information by moving around the central high-resolution area (*fovea*) through saccadic eye movements. But jumping spiders' eyes cannot move in their sockets.⁵ Thus, unlike mammals, they can't just move their eyes to focus on something at the periphery of their visual field. They have to move their whole body. Thus, the various movements by means of which jumping spiders turn to face, and track, salient environmental contingencies are functionally analogous to our saccadic eye movements - they're whole-body saccades (eg. Land 1969). So, if one is persuaded that information-gathering saccadic eye movements extend the human mind (see references above) one seems forced to concede that jumping spiders' minds are similarly extended.

⁵ They can however move the retinae of their anterior median eyes to selectively "bring into focus" various bits of the perceptual image they capture.

Further, in at least some species of jumping spiders these whole-body saccades play further cognitive roles. *Portia fimbriata* is one such species. These spiders boast an impressive, even “mammalian level” ability to *detour* (see Tarsitano & Jackson 1994; 1997). Detouring is the ability to reach one’s prey through very *indirect* paths, even when such paths force one to lose sight of the prey, or to temporarily move away from it. But how are these detours planned? *Portia* brains seem just too little to have the horsepower needed to create and update a spatial map of the spider’s surroundings. And in fact, *portia* do not plan these detours using an internal map. Rather, complex series of whole-body saccades allow these spiders to *act out* the evaluation process whereby the detour is selected via a process of vicarious trial and error. Vicarious trial and error is typically taken to be the behavioral manifestation of an inner deliberative process, through which various options are assessed so as to pick the best one (eg. Reddish 2016). In the case of *portia*, however, the “behavioral manifestation” seems to be part and parcel of the deliberative process. Simplifying, spiders orient themselves towards their goal, identify horizontal features (that is, pathways) departing from it, and trace the length of such features through whole-body saccades. When the line breaks, the spider reorients towards the goal and tracks another feature. Through repeated iterations of this procedure, the spider traces and “tests” numerous paths (horizontal lines) bringing it to the goal, eventually selecting the one it has been looking at the most, as that is the most uninterrupted path leading to the goal (Tarsitano & Andrews 1998; Tarsitano 2006).

In this latter case, one could argue in favor of extended spider cognition also on a non-analogical basis. In fact, a straightforward application of the parity principle (eg. Clark 2008, p. 77) guarantees the extension. The parity principle is a rule-of-thumb test to adjudicate cases of cognitive extension. When considering a putative piece of extended cognitive processing, the parity principle asks us how we would evaluate a *cerebral* functional isomorph of that piece of processing. If we would deem the latter as cognitive, then we need to recognize the former as cognitive too; else, we would fall prey to a neurocentric bias (for further discussion, see Clark 2011; Wheeler 2011, 2019b). Now, given that, when neurally realized, vicarious trial and error *is in fact* typically considered cognitive in both animals and humans (eg. Tolman 1939), one can safely conclude that *portia*’s extended realization of vicarious trial and error is a case of extended cognitive processing.

Let us state explicitly, for the sake of clarity, that we are *not* claiming that the physical machinery whereby jumping spiders perform whole body saccades is extended. Whilst that machinery is in a sense extended - after all, the physical machinery producing these movements will include the actuators causing the limbs to move, which are extra-cerebral - that is *not* the point we are making. The point we are making is that *the whole body saccades themselves* are part of the jumping spiders’ cognitive processing, just like saccadic eye movements (see O’regan 1992; Clark 2008, pp. 11-15; Rowlands 2010) and whatever cerebral process responsible for vicarious trial and error are part of the cognitive processing of humans (and rats).

In conclusion, jumping spiders’ whole-body saccades provide us a potent and empirically well supported case in favor of the extension of animal minds. Notice that, even if such an example were to be defused, accepting that human saccadic eye movements extend human cognition likely entails that animal cognition is extended too. For, most likely, the saccadic eye movements of other primates - and mammals more generally - will play similar functional roles in their cognitive lives. But, if these roles are cognitive in the human case, then they will most likely be cognitive in the animal case too.

That should not be surprising. For, “prop-less” epistemic actions seem to be widely distributed in nature; so much so, that one might find them performed even by what *might*⁶ be the most basic cognitive systems around, namely unicellular organisms. The bacterium *Salmonella Typhimurium*, for example, determines in which way to swim by comparing whether the concentration of nutrients “sampled” at two different times

⁶ Recall that, for the purpose of this paper, we are officially neutral on whether unicellular organisms qualify as cognizers or not.

has increased. But in order for such a comparison to make sense, the bacterium must constantly be on the move, so as to sample different points of its environment - otherwise, it will always be unable to detect any nutrient gradient (eg. Macnab & Koshland 1972)! So, *if* such bacteria are cognizers, it could be appropriate to depict them as active information foragers, locked in a perennial series of cognition-extending epistemic actions.

To sum up, animals are capable of executing “prop-less” epistemic actions. Thus, just like human cognition, animal cognition seems able to breach the sensorimotor frontier, leaking into the environment thanks to the purposeful motion of their active bodies.

A defender of an anthropocentric reading of the extended mind, however, might argue that human and animal extended cognitive systems are still worlds apart. Human cognition extends *also* by means of external information-encoding props such as tools and external representations. These external props permanently alter the human cognitive niche, thereby allowing for novel cognitive extensions to build on top of older ones. In this way, human cognitive extension improves through time, allowing humans to reach otherwise impossible cognitive achievements (eg. Menary 2007, 2015; Wheeler and Clark 2008, Malafouris 2013; Fabry 2020). Animal cognitive extensions, in contrast, are impermanent epistemic actions.

As we will soon argue, however, animal cognition extends via external, information-encoding props too.

4 - Cognition-extending props in animals

Human cognitive extensions often involve external, information-encoding props - consider, for instance, the examples of cognitive extension given in §1. Can animal cognitive extensions involve similar props? We’ll argue that yes, they can. But two *caveats* are needed beforehand.

First, we will interpret “external props” in a very broad way, so as to include all sorts of tools, artifacts (i.e. objects or structures *built* to achieve a goal) and naturefacts (i.e. naturally occurring objects or structures that allow us to achieve a goal, like a rock used to hammer a nail, eg. Hilpinen 2011; Preston 2022). This seems to be in line with the range of possible cognition-extending props in the human case, which range from complex technologies to very mundane things like a stick planted in the sand to create a primitive sundial.⁷ Secondly, our aim here is not that of establishing that animals use tools - for that has already been established on solid empirical grounds (eg. Bentley-Condit *et al* 2010). Our purpose here is to show that animals interact with certain relevant props in a mind-extending way. So, not *every* case of animal prop usage will do.

One paradigmatic case of prop-involving human cognitive extension is *cognitive offloading* (eg. Risiko & Gilbert 2016). During cognitive offloading, an agent delegates part of the cognitive processing to the environment, using external cognitive resources *rather than* internal ones. For instance, by writing a shopping list, one “delegates” the remembering to list itself, using it *rather than* an internally encoded memory representation. Many different animals seem to offload cognition when building their nests. Consider, for example, mud wasp of the genus *paralastor*. These wasps use mud to build fairly complex nests, consisting of an inner chamber (buried quite deep) and a fairly long entrance that protrudes several centimeters above the ground. Crucially, the entrance is not just a straight tube. Rather, its entrance is bent downwards, and expands so as to assume a bell-shaped form. How do *paralastor* wasps manage to build such nests? One option is that they build their nests consulting an inner model or set of instructions that “tells”

⁷ Of course, we don’t wish to claim that *any* prop can become a cognitive extension. For example, in order to be used as a cognitive extension, a prop must be deployable in a *transparent* manner; that is, in a way such that the deployment of the prop is swift, irreflexive and automatic, and the prop’s usage does not pose any problem taxing the agent’s cognitive resources (see Facchin 2022; Smart *et al.* 2022 for up-to-date analysis of transparency in EC).

them what to do. But this internal-processing heavy strategy does not seem the one adopted. Rather, they use the nest *itself* as the relevant instructions to follow. For, *paralastor* wasps seem to be exquisitely sensitive to certain perceptual cues, which immediately trigger certain nest building behaviors. For example, seeing the entrance tube protruding over a certain, critical, threshold will cause the wasp to shift behavioral patterns: she will stop elongating the tube and start bending it to then work at the bell-shaped entrance. But if mud is then added to the base of the nest, so that the entrance tube does not reach the threshold level, the wasp will simply go back to her elongating behavior - as if oblivious of the overall shape of the nest. Similarly, the perception of a hole in the mud causes the wasp to start building an entrance tube - even when the hole has been carved *on top* of the bell-shaped entrance to the nest (see Smith 1978).

So, it appears that *paralastor sp.* does not harbor an inner model or set of instructions “telling” them how their nests should be built. The triggering of appropriate actions has been delegated to the relevant perceptual features of the model, which are used “instead of” any inner instruction representing how to build the nest. If cognitive offloading extends our human minds, then, it appears that wasps’ minds are extended too. And not just wasps’ minds that are extended. For, a similar offloading strategy has been observed in nest-building spiders (for a summary focused on cognitive extensions, see Japyassú & Laland 2017) and even in certain species of birds (see Collias and Collias 1962).

Furthermore, this kind of self-cueing has been recognized as a special case of *stigmergy*, that is, the usage of signals to indirectly manipulate and organize an agent’s behavior (see Heyligen 2015, 2016a,b). This is important to notice for two reasons. First, it tells us that the relevant cues are signals: they carry *information* about what to do, and are thus likely used as *representations* of what to do. This is important to notice, for if the relevant cues are *used as* representations, then the agent behavior is not just a blind, mechanical and reflexive tropism triggered by a mere stimulus. For, an agent sensitive to the content of representations is likely behaving in a mindful, cognition-revealing, manner.⁸ Secondly, since stigmergic behaviors involve dense and bidirectional causal loop that are markedly non-linear (see references above), noticing that the case of cognitive offloading thus far discussed are case of stigmergic behavior opens up the space for a *second* argument to the effect that these animals’ cognitive processing is extended. In fact, several authors take such loopy non linear dynamics to support the extension of cognition in the human case (Lamb and Chemero 2018). We just see no reason as to why they *shouldn’t* support the extension of cognition in the animal case. Thus, they offer a second argument for our conditional claim: if loopy dynamics extend cognition in the human case, they also extend it in the animal case.

One might worry that the cases thus far examined do not suffice to substantiate our claim. For, in standard cases of human cognitive offloading, the *epistemic* actions whereby one offloads information and then deploys it are clearly separated from the *pragmatic* actions whereby one attains one’s goal. When one writes and consults a shopping list, one’s pragmatic goal is not to read/write something, but rather to *buy* certain things. The writing and reading are not goals *per se*, they’re only brought about to access information. But in the case examined above no such neat separation between epistemic and pragmatic action holds. The relevant perceptual cues are brought about and used when the animal is pursuing a *pragmatic* goal and are directly involved in its attainment. Thus, they might not qualify as epistemic actions in any relevant sense - thereby failing to extend animal minds.

⁸ One could object that in at least some cases the “tropism-based” reading should be preferred. After all, aren’t mud wasps the *go-to case* of an animal whose behavior is orchestrated by cunningly arranged tropisms? Our answer is “yes, but actually no”. Yes: mud wasps *are* the go-to case of behavior guided by tropisms. Yet, that traditional description is grossly exaggerated and misrepresents their behavior, which is not nearly as rigid and “mechanical” as it is traditionally described (see Keijzer 2001). So: actually no, we *shouldn’t* actually prefer the tropism based reading.

We believe such a worry is misguided for two reasons. First, there are paradigmatic cases of human cognitive offloading in which no such separation holds. Secondly, there are cases of animal cognitive offloading in which the relevant separation holds. Thus, the separation might not be needed; but even if it were, that would not damage our claim.

Consider a paradigmatic case of cognitive offloading: the usage of pen and paper (and numerals) as an external memory, to “remember” the partial results involved in a long mathematical operation (eg. Wilson 1994, Menary 2015). Now, imagine a school child dealing with a math test. Suppose her grade will not depend only on her final answer to the questions, but also by the procedures used to compute them. On the reasonable assumption that she’s aiming for a good grade, her writing down the relevant partial results and operations carried out is directly involved in bringing about the desired good. The barrier separating epistemic and pragmatic actions is thus breached, just as it was breached in the animal cases above. So, if in this case the child’s cognition extends (as many supporters of EC would argue), then no rigid distinction between pragmatic and epistemic actions is needed, and the worry above is defanged.

But what if the distinction were needed? Strange as it may sound one could technically hold that the teacher’s intentions in the example above prevent any cognitive extension from obtaining. Still, there would be cases of “purely epistemic” animal cognitive offloading. Consider, as an example, *patch marking*. Patch marking is a behavior found in numerous species of parasitoid insects, whereby the insect “marks”, often with a chemical marker, a previously explored patch of the environment. Importantly, these insects tend to respond aversively to these marks, spending less time in the “marked” regions. In this way, they do not have to recall *which* region of the environment they have already visited, the mark does the remembering for them (see Holler & Hormann 1993; Sheelan *et al.* 1993 Bernstein & Driessen 1996; Nakashima *et al.* 2002). Now, depositing the marks does not directly contribute to these insects’ goals (typically, injecting their eggs in a suitable host). So, patch marking seems an epistemic action separated from pragmatic ones in the desired way.

One could further object that, whilst in the animal cases here considered the relevant cognitive processes can’t be but extended, in the analog human cases the relevant cognitive process can take place in a purely intracranial way. For example, whilst we typically *do* use pen and paper to do math, we strictly speaking *can* do math purely in the head. Such a difference, the objector suggests, invalidates our analogy.

Yet, it is not clear that there is such a disanalogy, not that, if present, it would have the required significance. Can we really perform complex mathematical operations purely “in the head”? We think the answer is negative. Few people (if any) would be able to add 748948495947399359484762349032 to 238394038728972856458264 in the head. It seems our brains simply lack the working memory to carry out such an operation - which is just a simple addition. And indeed, our ability to perform even basic arithmetical operations seems to be an enculturated cognitive capacity, whose exertion depends *essentially* on our ability to manipulate external symbols according to rules (Fabry 2020). But even if this weren’t the case, and we really were able to perform complex arithmetic operations entirely “in the head”, it would still be true that, more often than not, we *don’t*, and rely on pen and paper (or calculators). So, as a matter of fact, most times our arithmetical cognition *is* extended. And, once that is conceded, our analogy delivers the desired conclusion.

Summarizing: several animals seem able to manipulate their own environments so as to offload certain cognitive processes to the environment itself. So, if cognitive offloading extends the cognitive systems of humans, then the cognitive systems of certain animals will be extended too.

5 - Allaying some worries

Thus far, we have argued that *if* human cognition extends, then animal cognition extends too. Human and animal cognitive systems are either *both* extended or *both* non-extended. Either way, cognitive extensions are not a human-only affair, nor do they tell humans apart from all other animals. EC should thus lose its anthropocentric focus. We know such a claim might generate some worries. Here, we allay the ones we find the most pressing.

First, one may worry our claim rests on a conceptual sleight of hand. Our discussion started in §1 with intrabodily cognitive extensions, which are tied to the “first wave”, computational and representational rendition of EC. But then, subsequent sections veered towards a more sensorimotor, dynamical and enactive rendition of EC, based on loopy and non-linear interaction dynamics. That’s the trademark of “second” and “third wave” EC (eg. Gallagher 2018; Wheeler 2019b on the “weaves” of EC) or even a symptom that we’ve been relying on a stronger thesis, according to which cognitive systems are *always* constitutively extended, almost by definition (Hutto *et al* 2014). Such a conceptual shift invites three problems. First, if cognition is always, “by definition” extended, then the claim that animal cognition is extended becomes trivial. Secondly, these renditions of EC hinge on some contested assumptions - for example, to anti-representationalism - that are not justified in our paper. Lastly, the claim that *all* agent-environment interactions extend an agent cognition (eg. Hutto *et al.* 2014) seems to lead to a *reductio* of EC. Surely defenders of EC should allow that an agent might find it useful to go for a walk to think without thereby having to count *roads* as part of the agent’s extended cognitive system!

These are genuine problems. But they do not affect us here. For, we’ve carefully avoided committing to *any* particular variant of EC. And we don’t need to commit. Our claim is that if certain props or processes extend the human mind *somehow*, then they also extend some animal minds. Of course, to maintain this agnosticism we had to discuss a number of examples compatible with various renditions of EC. Thus, the perceived shift does not reveal an implicit commitment to contested claims concerning the non-representational, always extended nature of cognition. Which, to be crystal clear, does not entail that we’re here accepting the claims concerning the representational nature of cognition and suggesting a sort of cognitive primacy of the brain made by “first-wave EC”.⁹ We remain neutral on such matters.

For the same reason, we’re not moving *too* fast by not discussing all the various objections to EC, as the *second worry* goes. For, all standard objections EC hinge on *some* assumption made only by *some* rendition of EC, but not others. For instance, objections based on representational contents (Adams and Aizawa 2001) are pretty toothless against non-representational renditions of EC (eg. Chemero 2009, Hutto & Myin 2013, pp. 139-145). Since here we’re not committed to *any* specific rendition of EC, we do not feel the need to engage with *any* of these objections. Moreover, our claim has a *conditional* form: if EC (in *any variant*) is true of humans, then EC is true of at least some animals too. Since the objections to EC typically deny the antecedent of this conditional, they don’t pose any threat to the truth of our claim.

A third, related, worry is that the cases we have been examining in §§ 2-4 do not meet certain criteria necessary for cognitive extension. This worry basically aims at limiting the validity of our conditional claim. It contends that it is not true that if human cognition extends *then* animal cognition extends too - in order to be extended, animal cognition must also meet certain specific conditions that separate genuine cases of cognitive extension from mere (i.e. not cognition extending) environmental embedding.

⁹ Whether or not “first wave EC” *actually* makes such claims is a complex matter we won’t discuss here (see Clark 1998, 2011 for discussion).

This *third worry* is onto something: surely we don't wish to claim that every animal-environment interaction extends the animal's cognition. Nor do we wish to claim that the existence of human cognitive extensions entails the existence of animal cognitive extensions. Such an entailment holds only *ceteris paribus*, that is, supposing that the putative animal extension meets the criteria for being an extension. Sadly, however, there is no consensus on what the relevant criteria to separate genuine extensions from "mere" environmental embedding are. Still, we wish to argue that the cases we discussed meet at least one, very popular, set of such criteria, namely the so-called "trust and glue" criteria (see Clark & Chalmers 1998). These criteria are: (i) that a resource extending one agent's cognition is readily available and typically invoked, (ii) that the information retrieved by relying on a resource is automatically endorsed (i.e. not consciously scrutinized before use) and (iii) that the relevant information retrieved by the resource is readily available when needed.¹⁰

It is straightforward to show that the extensions we discussed above satisfy (i)-(iii). The head movements of birds, the "whole body saccades" of spiders, as well as the various types of cues and mnemonic traces that various species of wasps use in their activities are all easily performed or produced. Indeed, as far as we can see these animals never omit to perform or produce them when facing the relevant cognitive tasks. Hence, the movements or cues that in our view extend these animals' cognition are readily available and typically invoked, as required by (i). As far as we know, moreover, none of these animals has ever been observed to stop and ponder on the information these movements or cues deliver. No bird, for example, has been observed to metacognitively evaluate the information gained through head-bodding, nor has any parasitoid wasp ever stopped to ponder on the reliability of the patch marking they use. It seems that these animals simply trust the relevant information and deploy it in the relevant contexts. Thus, the information seems to be automatically endorsed, as required by (ii). And given the ease with which the relevant movement can be performed and the relevant cue can be produced, the relevant information seems to be easily available when needed. Hence, (iii) is satisfied too. An analogous line of reasoning holds for intrabodily cognitive extensions (see Facchin *et al.* 2021, pp. 10-11; Boem *et al.* 2021, pp. 21-23). Thus, the cases that we've been discussing here do qualify as genuine cases of cognitive extension, at least according to the popular "trust and glue" criteria.

Importantly, the same seems to hold true even if one ditches "trust and glue" criteria for other sets of criteria. This is important to notice, as it signals that the cases we've been discussing qualify as cases of cognitive extension in a reliable manner, even when the relevant criteria for extendedness are changed. This suggests that they qualify as cognitive extensions not because of some "blindspot" or idiosyncratic feature of the trust and glue criteria, but because these cases are genuine cases of extended cognition. For example, some philosophers have argued that "trust and glue" criteria should be substituted by "mutual manipulability" criteria (see Kaplan 2012; Krickel 2019).¹¹ Very roughly put, the idea is that a putative cognitive extension qualifies as such if the putative extension and the cognitive phenomenon or process it is said to extend are mutually manipulable; that is, if manipulations of the putative extension alter the phenomenon or process, and *vice versa*. It is easy to see that the cases we described in §§ 2-4 meet this criterion. Just to give an example: altering a jumping spider's pattern of "whole-body saccades" (i.e. a manipulation of the putative cognitive extension) will alter the spider's route selection process - since the spider selects the route it has been looking at the most, then making it look at a certain route more than other will impact the route it chooses. And *vice versa*: manipulating the spider's selection process (e.g. making it choose amongst different

¹⁰ Originally, Clark & Chalmers (1994) included a fourth condition concerning the previous conscious endorsement of the relevant information, but the status of this fourth criterion has always been contested, and it is typically expunged from the "trust and glue" criteria (eg. Clark 2008).

¹¹ Even if Beate Krickel might have recently changed her mind, see (Krickel 2023).

routes) will make the spider act out different “saccadic whole-body” movements (as detailed in the examples discussed in §3).

A *fourth worry* is that our claim is based on the implicit adoption of a “biogenic” conception of cognition, according to which pretty much every adaptive, flexible and information-guided organism-environment interaction qualifies as an instance of *full-blown* cognition (see Lyon 2006, 2020). Many find such a conception of cognition objectionable, as it fails to properly identify cognitive processes, states and agents (see Adams 2018). As such, “biogenic” cognition should be rejected, together with our claim.

However, as indicated in §1, we do not subscribe to any particular conception of cognition. Indeed, as claimed there, we do identify cognition by looking at paradigmatic cases - and this is just not how the biogenic approach identifies cognitive processes and phenomena. According to the biogenic approach, for example, epigenetic memory is a *genuinely cognitive* case of memory, in the exact same sense semantic memory and working memory are (e.g. Lyon 2015). Our view, however, does not identify epigenetic memory as a cognitive type of memory - for it is not a paradigmatic example of cognition (nor of memory, for that matter). As such, we leave epigenetic memory (and several other similar cases) in a “limbo” of sorts, in which they are neither deemed cognitive nor not-cognitive. We recognize that leaving such a “gray area” may be disappointing for some, but we feel that this is the best that can be realistically done given the absence of any non-partisan and widely shared conception of cognition (see references given in §1).¹²

In *fifth place*, one may worry that, if true, our claim would merely replace a form of human cognitive exceptionalism with a form of *animal*¹³ cognitive exceptionalism. What we’ve shown, this worry goes, is that only animals’ cognitive systems are extended. But what about other cognitive agents such as plants (eg. Calvo & Keijzer 2011), robots (eg. Tani 2016), single-celled organisms (eg. Lyon 2015) and so-called “intelligent” materials (eg. Mcgivers 2019)? By being silent on those, we reinforce a prejudiced view of cognition that sees it as an animal-only affair (eg. Calvo 2022).

This worry misconstrues our argumentative strategy. As we said above (§1) our focus on animals is partly instrumental: it allows us to streamline our argument without having to debate on whether systems such as plants, fungi and bacteria are *really* cognitive. Importantly, were such systems cognitive, nothing of what we’ve argued thus far would prevent their cognition from extending. So, our claim does not support a form of animal cognitive exceptionalism.

Conversely, the *sixth worry* is that our anti-anthropocentric argument is far too glib. After all there still are relevant differences between human and animal extended cognitive systems. Parasitoid insects might offload some small piece of their cognitive processing - but that’s not the same as using pen and paper to do a complex mathematical operation. Human cognitive extensions allow to individually improve on them and innovate (eg. Wheeler and Clark 2008). Further, human cognitive extension often require the agent to master some cultural practices with a relevant normative dimension: only certain ways of manipulating numerals are *right* (eg. Menary 2007, 2015; Fabry 2020). They’re also highly personalized and the information they convey will be trusted in a special way (eg. Sterelny 2010). Insects’ extensions, in contrast, are most likely genetically determined, in a way that does not allow for extensions to be individually improved and personalized, and does not require the agent to learn any cultural practice. These, the worry goes, are *significant* differences our claim simply fails to acknowledge.

¹² For the same reasons, we are not claiming that the cases examined in §§ 2-4 are cases of cognitive processes *because* they are described as cognitive by the researchers reporting on them. We are claiming that the processes in such cases are cognitive because they are instances of paradigmatically cognitive processes.

¹³ Here intended in the proper biological sense, inclusive of humans .

These are significant differences indeed - but our claim is compatible with their acknowledgement. For the view that cognitive extension is not a human-only affair can logically co-exist with the view that human extensions and animal extensions have significant differences. Compare: it is typically accepted that tool-use is *not* exclusively human. But it is also typically accepted that human and animal tool use are importantly different (eg. Heersmink 2022). Does this mean that we're forced to recognize a difference *in kind* between human and animal's cognitive extensions? We're not sure. Whilst comparing human cognitive offloading and insects' cognitive offloading makes it tempting to give a positive answer, things look much murkier when it comes to human *vs* animals' epistemic actions (§4). Further, there seems no significant difference when it comes to intrabodily cognitive extensions. Perhaps, then, to correctly answer this question we might have to dissolve the generic kind "cognitive extension" into more specific sub-types of cognitive extensions, just as the generic kind "memory" has been dissolved into various sub-types of memories like episodic memory and semantic memory.¹⁴ Be as it might, our claim establishes that if there are such things as cognitive extensions, then they are widespread in the animal kingdom. As such, cognitive extensions are not uniquely human, and thus, cognitive extensions alone cannot be what makes humans unique.

6 - Conclusion

This paper began (§1) by noticing that EC often neglects animal minds, thereby acquiring an unwarranted human-centric focus. We've argued that such a focus ought to be abandoned, for, if human cognition extends, then animal cognition extends too. In fact, many animals have the same intrabodily cognitive extensions humans have (§2), perform epistemic actions comparable to ours (§3) and are even capable of offloading cognition in the environment through the usage of specific environmental props (§4). We then defended our claim from a number of objections (§5).

In closing, we wish to quickly point out an implication of our claim. Consider a case reported in (Clark 2003, pp. 139-141; 2023, pp.151-154). Patients suffering from Alzheimer's syndrome were tested, obtaining incredibly low scores in standard psychological tests. Thus, according to the tests, the patients were unable to live alone and take care of themselves. And yet they did. How? By cleverly structuring their environments, so as to offload on it the cognitive load their brains were no longer able to handle. Psychological testing, then, failed to take into account the entirety of these agent's cognitive systems - it tested only a part of it. And indeed, to separate these agents from their environment would amount to amputating a part of their cognitive system - at least if the EC story is correct.

If our claim is right, however, we might have been making the same mistake when it comes to the cognition of animals - as animals are often tested in artificial, un-ecological settings (eg. Barrett 2016). Clearly, these controlled, neat, and often quite schematically organized settings are crucial to control for the largest number possible of intervening variables. Ecologically realistic settings - qua ecologically realistic settings - would contain large amounts of noise, confounds, and unforeseen circumstances and events that would clearly mar one's experimental design. So, we are not suggesting ethologists should burn their labs down and go work in the forest. Ecologically unrealistic lab work is clearly essential. But, if our arguments are correct, the result it delivers should be taken with a pinch of salt (or at least with a pinch of salt more than the pinches of salt already used to interpret these data). For, just like the patients in Clark's case, we might be inadvertently amputating the cognitive systems of animals just by testing their cognitive and behavioral capabilities in the lab.

References

Adams, F. (2018). Cognition wars. *Studies in History and Philosophy of Science Part A*, 68, 20-30.

¹⁴ On dissolving kinds see (Ramsey 2021).

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical psychology*, 14(1), 43-64.
- Akagi, M. (2018). Rethinking the problem of cognition. *Synthese*, 195(8), 3547-3570.
- Allen, C., et al. (2009). The lower bounds of cognition: what do spinal cords reveal?. In Bickle, J. (Ed.). *The Oxford Handbook of Philosophy and Neuroscience* (pp. 129-142). New York: Oxford University Press
- Allen, C. (2017). On not defining cognition. *Synthese*, 194(11), 4233-4249.
- Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human-machine interface. *Trends in cognitive sciences*, 7(12), 541-546.
- Barrett, L. (2016). Why brains are not computers, why behaviorism is not satanism, and why dolphins are not aquatic apes. *The Behavior Analyst*, 39(1), 9-23.
- Bentley-Condit, V. et al. (2010). Animal tool use: current definitions and an updated comprehensive catalog. *Behaviour*, 147(2), 185-32A.
- Bercik, P., et al. (2011). The intestinal microbiota affect central levels of brain-derived neurotropic factor and behavior in mice. *Gastroenterology*, 141(2), 599-609.
- Bocanegra, B. R., et al (2019). Intelligent problem-solvers externalize cognitive operations. *Nature Human Behaviour*, 3(2), 136-142.
- Boem, F., et al (2021). Out of our skull, in our skin: the Microbiota-Gut-Brain axis and the Extended Cognition Thesis. *Biology & philosophy*, 36(2), 1-32.
- Boem, F., et al. (2024). Minding the gut: extending embodied cognition and perception to the gut complex. *Frontiers in Neuroscience*, 17, 1172783.
- Bernstein, C., & Driessen, G. (1996). Patch-marking and optimal search patterns in the parasitoid *Venturia canescens*. *Journal of Animal Ecology*, 211-219.
- Calvo, P. (2022). *Planta Sapiens*. Little Brown.
- Calvo, P., & Keijzer, F. (2011). Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive behavior*, 19(3), 155-171.
- Carter, A., et al. (2018). Introduction. In A. Carter, A. Clark, J. Kallestrup, O. Palermos, D. Pritchard (eds.) *Extended Epistemology* (pp. 1-16). Oxford University Press.
- Cao, R. (2022). Multiple realizability and the spirit of functionalism. *Synthese*, 200: 506.
- Chalmers, D. (2008). Foreword. In A. Clark., *Supersizing the Mind*. New York: Oxford University Press.
- Chalmers D. (2019). Extended cognition and extended consciousness. In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 9-21). New York, Oxford University Press
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. Cambridge, MA.: The MIT Press.
- Clark, A. (1997). *Being There*. Cambridge, MA.: The MIT Press.
- Clark, A. 1998. "Author's Response." In "Review Symposium on Andy Clark's Being There." *Metascience* 7: 95-103

- Clark, A. (2003). *Natural-Born Cyborgs*, New York: Oxford University Press.
- Clark, A., (2008). *Supersizing the Mind*. New York: Oxford University Press.
- Clark, A. (2011). Finding the mind. *Philosophical Studies*, 152(3), 447-461.
- Clark, A. (2019). Replies to Critics, In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 266-300). New York, Oxford University Press.
- Clark, A. (2022). Extending the Predictive Mind. *Australasian Journal of Philosophy*, <https://doi.org/10.1080/00048402.2022.2122523>
- Clark, A. (2023). *The experience machine*. London: Penguin.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7-19.
- Collias, N. E., & Collias, E. C. (1962). An experimental study of the mechanisms of nest building in a weaverbird. *The Auk*, 79(4), 568-595.
- Colombo, M., Irvine, L., & Stapleton, M. (2019). *Andy Clark and His Critics*. New York: Oxford University Press.
- De Cruz, H., & De Smedt, J. (2013). Mathematical symbols as epistemic actions. *Synthese*, 190(1), 3-19.
- Dutilh-Novaes, C. (2013). *Formal Languages in Logic*. New York: Oxford University Press.
- Drayson, Z. (2010). Extended cognition and the metaphysics of mind. *Cognitive Systems Research*, 11(4), 367-377.
- Fabry, R. E. (2020). The cerebral, extra-cerebral bodily, and socio-cultural dimensions of enculturated arithmetical cognition. *Synthese*, 197(9), 3685-3720.
- Facchin, M. (2022a). Why can't we say what cognition is (at least for the time being). Preprint. <http://philsci-archive.pitt.edu/20348/>
- Facchin, M. (2022b). Phenomenal transparency, cognitive extension, and predictive processing. *Phenomenology and the Cognitive Sciences*, 1-23.
- Facchin, M., et al. (2021). Retiring the "Cinderella view": the spinal cord as an intrabodily cognitive extension. *Biology & Philosophy*, 36(5), 1-25.
- Favela, L. H., & Martin, J. (2017). "Cognition" and dynamical cognitive science. *Minds and Machines*, 27(2), 331-355.
- Figdor, C. (2018). *Pieces of Mind*. New York: Oxford University Press.
- Figdor, C. (2021). The psychological speciesism of humanism. *Philosophical Studies*, 178(5), 1545-1569.
- Figdor, C. (2022). What could cognition be, if not human cognition? *Biology and Philosophy*, 37(52), 1-21.
- Figdor, C. (2023). What Are We Talking About When We Talk About Cognition?: Human, Cybernetic, and Phylogenetic Conceptual Schemes. *Jolma*, 4(2), 149-162.

- Fiori F., et al. (2014). Motor imagery in spinal cord injury patients: moving makes the difference. *Journal of Neuropsychology*, 8(2), 199-215
- Foster, J. A., & Neufeld, K. A. M. (2013). Gut–brain axis: how the microbiome influences anxiety and depression. *Trends in neurosciences*, 36(5), 305-312.
- Fujita M. (2006). Head-bobbing and non-bobbing walking of black-headed gulls (*Larus ridibundus*). *Journal of comparative physiology. A, Neuroethology, sensory, neural, and behavioral physiology*, 192(5), 481–488.
- Fux, M., & Eilam, D. (2009). How barn owls (*Tyto alba*) visually follow moving voles (*Microtus socialis*) before attacking them. *Physiology & behavior*, 98(3), 359–366.
- Gallagher, S. (2018). The extended mind: state of the question. *The Southern Journal of Philosophy*, 56(4), 421-447.
- Gareau, M. G., et al. (2011). Bacterial infection causes stress-induced memory dysfunction in mice. *Gut*, 60(3), 307-317.
- Gough, J. (2023a). Cognitive science meets the mark of the cognitive: putting the horse before the cart. *Biology & Philosophy*, 38(1), 1.
- Gough, J. (2023b). Between mind and body? Psychoneuroimmunology, psychology, and cognitive science. *Perspectives on Science*, 1-38.
- Grau, J. W. (2014). Learning from the spinal cord: how the study of spinal cord plasticity informs our view of learning. *Neurobiology of learning and memory*, 108, 155-171.
- Heersmink, R. (2022). Human uniqueness in using tools and artifacts: flexibility, variety, complexity. *Synthese*, 200(6), 1-22.
- Heylighen, F. (2015). *Stigmergy as a Universal Coordination Mechanism: components, varieties and applications. Human Stigmergy: Theoretical Developments and New Applications*; Springer: New York, NY, USA.
- Heylighen, F. (2016a). Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research*, 38, 4-13.
- Heylighen, F. (2016b). Stigmergy as a universal coordination mechanism II: Varieties and evolution. *Cognitive Systems Research*, 38, 50-59.
- Hilpinen, R. (2011). Artifact. In E. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/win2011/entries/artifact/>
- Höller, C., & Hörmann, R. (1993). Patch marking in the aphid hyperparasitoid, *Dendrocerus carpenteri*: the information contained in patch marks. *Oecologia*, 94(1), 128-134.
- Huber-Huber, C., Buonocore, A., & Melcher, D. (2021). The extrafoveal preview paradigm as a measure of predictive, active sampling in visual perception. *Journal of Vision*, 21(7), Article 12.
- Hurley, S. (2001). Perception and action: Alternative views. *Synthese*, 129, 3-40.

- Hutto, D. D., Kirchhoff, M. D., & Myin, E. (2014). Extensive enactivism: why keep it all in?. *Frontiers in human neuroscience*, 8, 706.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism*. Cambridge, MA.: The MIT Press.
- Japyassú, H. F., & Laland, K. N. (2017). Extended spider cognition. *Animal Cognition*, 20(3), 375-395.
- Jones, M.P., Pierce, K.E., & Ward, D.A. (2007). Avian vision: a review of form and function with special consideration to birds of prey. *Journal of Exotic Pet Medicine*, 16, 69-87.
- Kandel E., et al. (2012). *Principles of Neural Science* (fifth edition). London: The McGraw-Hill Companies.
- Kaplan, D. M. (2012). How to demarcate the boundaries of cognition. *Biology & Philosophy*, 27, 545-570.
- Keijzer, F. (2001). *Representation and behavior*. MIT Press.
- Krickel, B. (2019). Extended cognition, the new mechanists' mutual manipulability criterion, and the challenge of trivial extendedness. *Mind & Language*, 35(4), 539-561.
- Krickel, B. (2023). Extended Cognition and the Search for the Mark of Constitution—A Promising Strategy?. In *Situated Cognition Research: Methodological Foundations* (pp. 129-146). Cham: Springer International Publishing.
- Kirsh, D. (2019). When is a Mind Extended? In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 128-142). New York, Oxford University Press.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive science*, 18(4), 513-549.
- Kral K. (2003) Behavioural—analytical studies of the role of head movements in depth perception in insects, birds and mammals, *Behavioral Processes*, 64, 1, 1-12.
- Lamb, M., & Chemero, A. (2018). Interaction in the open: where dynamical systems become extended and embodied. In A. Newen, L. de Bruin, S. Gallagher (Eds.), *The Oxford Handbook of 4E Cognition*, (pp. 147-162). New York: Oxford University Press.
- Land, M. F. (1969). Movements of the retinae of jumping spiders (Salticidae: Dendryphantinae) in response to visual stimuli. *Journal of experimental biology*, 51(2), 471-493.
- Lisney T. J., & Troje N. (2016). Head-bobbing in the Ring-billed Gull (*Larus delawarensis*). *Canadian Field-Naturalist* 130(2): 174-177.
- Lyon, P. (2006). The biogenic approach to cognition. *Cognitive Processing*, 7(1), 11-29.
- Lyon, P. (2015). The cognitive cell: bacterial behavior reconsidered. *Frontiers in microbiology*, 6, 264.
- Lyon, P. (2020). Of what is “minimal cognition” the half-baked version?. *Adaptive Behavior*, 28(6), 407-424.
- Macnab, R., & Koshland, D. (1972). The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 69, 2509–2512.

- MacIver, M. A. (2009). Neuroethology: From morphological computation to planning. In Robbins, P., & Aydede, M. (Eds.). *The Cambridge handbook of situated cognition* (pp. 480-504). Cambridge University Press.
- Maglio, P., & Kirsh, D. (1992). Some epistemic benefits of action: Tetris, a case study. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (Vol. 14, p. 224).
- Malafouris, L (2013). *How Things Shape the Mind*. Cambridge, MA.: The MIT Press.
- McGivern, P. (2019). Active materials: minimal models of cognition? *Adaptive Behavior*, 28(6), 441-451.
- Meagher, M. et al. (1993). Activation of opioid and nonopioid hypoalgesic systems at the level of the brainstem and spinal cord: Does a coulometric relation predict the emergence or form of environmentally-induced hypoalgesia? *Behavioral Neuroscience* 107: 493–505.
- Menary, R. (2007). *Cognitive Integration*. Basingstoke: Pallgrave
- Menary, R. (2010). *The Extended Mind*. Cambridge, MA.: The MIT Press.
- Menary, R. (2015). Mathematical cognition: a case of enculturation. *Open Mind* 25(T), <https://doi.org/10.15502/9783958570818>
- Nakashima, Y., Teshiba, M., & Hirose, Y. (2002). Flexible use of patch marks in an insect predator: effect of sex, hunger state, and patch quality. *Ecological Entomology*, 27: 581- 587
- Neufeld, K. et al. (2011a). Reduced anxiety-like behavior and central neurochemical change in germ-free mice. *Neurogastroenterology & Motility*, 23(3), 255-e119.
- Neufeld, K. et al. (2011b). Effects of intestinal microbiota on anxiety-like behavior. *Communicative & integrative biology*, 4(4), 492-494.
- Nyakatura J. A., & Andrara E., (2014). On vision in birds: coordination of head-bobbing and gait stabilises vertical head position in quail. *Frontiers in Zoology*, 11, 27.
- O'regan, J. K. (1992). Solving the "real" mysteries of visual perception: the world as an outside memory. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 46(3), 461.
- Palermos, S. O. (2014). Loops, constitution, and cognitive extension. *Cognitive systems research*, 27, 25-41.
- Parise, A. G., Gagliano, M., & Souza, G. M. (2020). Extended cognition in plants: is it possible?. *Plant signaling & behavior*, 15(2), 1710661.
- Parise, A. G., et al. (2023) Ariadne's thread and the extension of cognition: a common but overlooked phenomenon in nature?. *Frontiers in Ecology and Evolution*, 10, 1304.
- Preston, B (2022). Artifact. In E. Zalta & U. Nodelman (Eds.) *The Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/win2022/entries/artifact/>
- Pritchard J. D., (2018) Situated Cognition and the Function of Behavior, *Comparative Cognition & Behavior*, 13, 35-40.
- Ramsey, W. M. (2021). What eliminative materialism isn't. *Synthese*, 199(3), 11707-11728.
- Reddish, A. d. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3), 147-159.

- Risko, E. F. & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688
- Rizzolatti, G., & Sinigaglia, C. (2006). *Mirrors in the Brain*. New York: Oxford University Press.
- Rowlands, M. (2010). *The New Mind Sciences*. Cambridge, MA.: The MIT Press.
- Scandola M., et al. (2016). Spinal cord lesions shrink peripersonal space around the feet, passive mobilization of the paraplegic limb restores it. *Scientific Reports*, 6, 24126.
- Scandola, M., et al. (2020). Visuo-motor and interoceptive influences on peripersonal space representation following spinal cord injury. *Scientific Reports*, 10(1), 1-16.
- Sedda A., et al. (2019). Affordances after spinal cord injury. *Journal of Neuropsychology*, 13(2), 354-369.
- Sheehan, W., Wäckers, F. L., & Lewis, W. J. (1993). Discrimination of previously searched, host-free sites by *Microplitis croceipes* (Hymenoptera: Braconidae). *Journal of Insect Behavior*, 6(3), 323-331.
- Sims, M., & Kiverstein, J. (2022). Externalized memory in slime mould and the extended (non-neuronal) mind. *Cognitive Systems Research*, 73, 26-35.
- Smart, P. R. (2021). Shedding Light on the Extended Mind: HoloLens, Holograms, and Internet-Extended Knowledge. *Frontiers in Psychology*, 12.
- Smart, P. R., Andrada, G., & Clowes, R. W. (2022). Phenomenal transparency and the extended mind. *Synthese*, 200(4), 335.
- Smith, A. P. (1978). An investigation of the mechanisms underlying nest construction in the mud wasp *Paralastor* sp. (Hymenoptera: Eumenidae). *Animal Behaviour*, 26, 232-240.
- Sterelny, K. (2003). *Thought in a Hostile World*. Wiley-Blackwell.
- Sterelny, K. (2010). Minds: extended or scaffolded?. *Phenomenology and the Cognitive Sciences*, 9(4), 465-481.
- Tani, J. (2016). *Exploring Robotic Minds*. New York: Oxford University Press.
- Tarsitano, M. (2006). Route selection by a jumping spider (*Portia labiata*) during the locomotory phase of a detour. *Animal Behaviour*, 72(6), 1437-1442.
- Tarsitano, M. S., & Andrew, R. (1999). Scanning and route selection in the jumping spider *Portia labiata*. *Animal Behaviour*, 58(2), 255-265.
- Tarsitano, M. S., & Jackson, R. R. (1994). Jumping spiders make predatory detours requiring movement away from prey. *Behaviour*, 131(1-2), 65-73.
- Tarsitano, M. S., & Jackson, R. R. (1997). Araneophagic jumping spiders discriminate between detour routes that do and do not lead to prey. *Animal Behaviour*, 53(2), 257-266.
- Theunissen L. M., & Troje N. F., (2017). Head Stabilization in the Pigeon: Role of Vision to Correct for Translational and Rotational Disturbances. *Frontiers in Neuroscience*, 11.
- Tolman, E. C. (1939). Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review*, 46(4), 318.

- Troje, N. F., & Frost, B. J. (2000). Head-bobbing in pigeons: how stable is the hold phase?. *The Journal of experimental biology*, 203(Pt 5), 935–940.
- Vahdat S., et al. (2015). Simultaneous brain-cervical cord fMRI reveals intrinsic spinal cord plasticity during motor sequence learning. *PLoS Biology*, 13(6). <https://doi.org/10.1371/journal.pbio.1002186>.
- Vallée-Tourangeau, F., et al (2016). Insight with hands and things. *Acta psychologica*, 170, 195-205.
- van den Brandeler, E. (2024). Towards an Epistemology of ‘Speciesist Ignorance’. *Res Publica*, 1-21.
- Varga, S. (2018). Demarcating the realm of cognition. *Journal for General Philosophy of Science*, 49, 435-450.
- Wang, T., et al (2015). *Lactobacillus fermentum* NS9 restores the antibiotic induced physiological and psychological abnormalities in rats. *Beneficial microbes*, 6(5), 707-717.
- Wheeler, M. (2011). In search of clarity about parity. *Philosophical Studies*, 152(3), 417-425.
- Wheeler, M. (2019a). The reappearing tool: transparency, smart technology, and the extended mind. *AI & society*, 34(4), 857-866.
- Wheeler, M. (2019b). Breaking the waves. In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 81-95). New York, Oxford University Press.
- Wheeler, M., & Clark, A. (2008). Culture, embodiment and genes: Unravelling the triple helix. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1509), 3563-3575.
- Wilson, R. A. (1994). Wide computationalism. *Mind*, 103(411), 351-372.
- Wilson, B. et al. (2015) Auditory sequence processing reveals evolutionarily conserved regions of frontal cortex in macaques and humans. *Nat. Commun.* 6, 8901.
- Wolpaw J. R., & Tennissen A. M. (2001). Activity-dependent spinal cord plasticity in health and disease. *Annual Review of Neuroscience*, 24(1), 807-843.